AD-A124 403 THE EFFECTS OF ITEM USAGE VARIATION ON INVENTORY STOCKAGE MODELS(U) AIR FORCE INST OF TECH WRIGHT-PATTERSON AFB OH SCHOOL OF SYST..

UNCLASSIFIED J H KUTZKE ET AL. SEP 82 AFIT-LSSR-77-82 F/G 15/5 NL

1/2

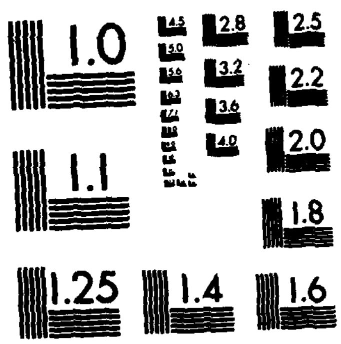
1.0
1.1
1.25
1.4
1.6
1.8
2.0
2.2
2.5
2.8
3.2
3.6
4.0

2

DEPARTMENT OF THE AIR FORCE

AIR UNIVERSITY (ATC)

AIR FORCE INSTITUTE OF TECHNOLOGY

Wright-Patterson Air Force Base, Ohio

THE EFFECTS OF ITEM USAGE VARIATION
ON INVENTORY STOCKAGE MODELS

Jonathan H. Kutzke, Captain, USAF
Gary A. Turner, GS-13

LSSR 77-82

The contents of the document are technically accurate, and no sensitive items, detrimental ideas, or deleterious information are contained therein. Furthermore, the views expressed in the document are those of the author(s) and do not necessarily reflect the views of the School of Systems and Logistics, the Air University, the Air Training Command, the United States Air Force, or the Department of Defense.

AFIT Control Number ~~LSSR 77-82~~

AFIT RESEARCH ASSESSMENT

The purpose of this questionnaire is to determine the potential for current and future applications of AFIT thesis research. Please return completed questionnaires to: AFIT/LSH, Wright-Patterson AFB, Ohio 45433.

1. Did this research contribute to a current Air Force project?

   a. Yes b. No

2. Do you believe this research topic is significant enough that it would have been researched (or contracted) by your organization or another agency if AFIT had not researched it?

   a. Yes b. No

3. The benefits of AFIT research can often be expressed by the equivalent value that your agency received by virtue of AFIT performing the research. Can you estimate what this research would have cost if it had been accomplished under contract or if it had been done in-house in terms of manpower and/or dollars?

   a. Man-years ________ $ ________ (Contract).

   b. Man-years ________ $ ________ (In-house).

4. Often it is not possible to attach equivalent dollar values to research, although the results of the research may, in fact, be important. Whether or not you were able to establish an equivalent value for this research (3 above), what is your estimate of its significance?

   a. Highly Significant b. Significant c. Slightly Significant d. Of No Significance

5. Comments:

________________________________ ________________________________

Name and Grade Position

________________________________ ________________________________

Organization Location

FOLD DOWN ON OUTSIDE - SEAL WITH TAPE

AFIT/LSH
WRIGHT-PATTERSON AFB OH 45433

OFFICIAL BUSINESS
PENALTY FOR PRIVATE USE, $300

NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES

BUSINESS REPLY MAIL
FIRST CLASS PERMIT NO. 73236 WASHINGTON D.C.

POSTAGE WILL BE PAID BY ADDRESSEE

AFIT/DAA
Wright-Patterson AFB OH 45433

FOLD IN

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER LSSR 77-82	2. GOVT ACCESSION NO. AD-A124403	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) THE EFFECTS OF ITEM USAGE VARIATION ON INVENTORY STOCKAGE MODELS		5. TYPE OF REPORT & PERIOD COVERED Master's Thesis
		6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s) Jonathan H. Kutzke, Captain, USAF Gary A. Turner, GS-13		8. CONTRACT OR GRANT NUMBER(s)
9. PERFORMING ORGANIZATION NAME AND ADDRESS School of Systems and Logistics Air Force Institute of Technology, WPAFB OH		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS
11. CONTROLLING OFFICE NAME AND ADDRESS Department of Communication and Humanities, AFIT/LSH, WPAFB OH 45433		12. REPORT DATE September 1982
		13. NUMBER OF PAGES 99
14. MONITORING AGENCY NAME & ADDRESS(if different from Controlling Office)		15. SECURITY CLASS. (of this report) UNCLASSIFIED
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report)

Approved for public release; distribution unlimited

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

18. SUPPLEMENTARY NOTES

APPROVED FOR PUBLIC RELEASE: IAW AFR 190-17

LYNN E. WOLAVER
Dean for Research and Professional Development

AIR FORCE INSTITUTE OF TECHNOLOGY (ATC)
WRIGHT-PATTERSON AFB, OH 45433

8 OCT 1982

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

Recoverable Models
Inventory Management
Variance-to-Mean Ratio
Demand Variation
Stockage Computation
Performance Calculation

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

Thesis Chairman: James M. Masters, Major, USAF

DD FORM 1 JAN 73 1473 EDITION OF 1 NOV 65 IS OBSOLETE

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

Although the USAF maintains accurate records of average usage data for recoverable items, very few historical data are available which record the variations in usage which items may exhibit. These usage variations or variances are necessary inputs to all USAF inventory models. Current USAF practice is not consistent with respect to how these variances are estimated and used in different stockage calculations, e.g., Peacetime Operating Stock (POS) requirements computations, distribution, WRSK/BLSS. The purpose of this research is to determine how possible errors in variance estimates will affect stockage models. A validated F-16 item usage data base and RAND's Dyna-METRIC stockage model were used to explore the effects on inventory investment costs and weapon system readiness and sustainability of errors in estimating item usage variances. It was found that both investment cost and readiness/sustainability were changed significantly when variances were changed. Implications are that the variance inputs to stockage models must be estimated accurately to avoid a misallocation of funds. Also, an accurate prediction of how a given stockage level will support readiness and sustainability depends on accurately determining the variance. This research suggests that the right variance-to-mean must be computed and used.

LSSR 77-82

THE EFFECTS OF ITEM USAGE VARIATION ON INVENTORY STOCKAGE MODELS

A Thesis

Presented to the Faculty of the School of Systems and Logistics

of the Air Force Institute of Technology

Air University

In Partial Fulfillment of the Requirement for the

Degree of Master of Science in Logistics Management

By

Jonathan H. Kutzke, BS
Captain, USAF

Gary A. Turner, BS
GS-13

September 1982

Approved for public release;
distribution unlimited

This thesis, written by

Captain Jonathan H. Kutzke

and

Mr. Gary A. Turner

has been accepted by the undersigned on behalf of the faculty of the School of Systems and Logistics in partial fulfillment of the requirements for the degree of

MASTER OF SCIENCE IN LOGISTICS MANAGEMENT

DATE: 29 September 1982

COMMITTEE CHAIRMAN

ACKNOWLEDGEMENTS

This thesis would not have been possible without those persons at HQ AFLC, Ogden Air Logistics Center, AFIT, and the RAND Corporation who so ably assisted us in our efforts. We extend our sincerest thanks to this group of true professionals.

Our wives, Vicki and Cynthy, and families deserve a special thanks. Many, many hours that should have been spent with them were spent with this thesis. Their tolerance and loving support inspired us to complete this research. Again, a heartfelt thanks.

We thank Suzanne for typing our figures and tables. She did a superb job.

And finally, to Vicki, our typist and editor, we extend our warmest and sincerest thanks. Our appreciation for the many, many hours she spent untiringly typing and retyping this document because of changes we had made, can never be adequately expressed with words. Her cheerful attitude was an inspiration to us both. Without her, this thesis simply wouldn't have gotten done.

TABLE OF CONTENTS

CHAPTER Page

CHAPTER	Page

LIST OF TABLES

LIST OF FIGURES

CHAPTER I

INTRODUCTION

Background

The United States Air Force exists to execute the defense policy of the United States. To be effective in this mission, the Air Force must maintain a state of readiness to go to war when called to do so. The readiness level or capability to wage war is dependent upon the proper levels and management of personnel, equipment (e.g., aircraft, starter carts, trucks, computers, etc.), and supplies (e.g., food, fuel, spare parts, etc.). It is clear that if the Air Force is to wage war quickly and effectively, it must manage its peacetime resources wisely and efficiently. Traditional peacetime thinking (and Congressional funding) led Air Force executives into concentrating on acquiring the needed weapon systems and maintaining personnel levels and spare parts to operate them in peacetime only. The nation could always mobilize for war if necessary. In today's nuclear age, however, with lead times of complex weapon systems and spare parts so long (years), and the devasting and rapid destructive capability of potential adversaries, there would be no time for mobilization of the industrial base. The "next

war" will most probably be fought quickly with existing personnel, weapons, and supplies.

The Air Force could have all the necessary personnel and equipment that it required to go to war, but if it lacked supplies to support those forces, then their effectiveness and sustainability would be severely jeopardized. The supply area can be broken into two basic groups: consumables and reparables. Munitions are a third group, but can be thought of as fitting more with consumables. The two basic groups are identifiable by the way they are used as well as the methods used for their management. Consumable items are those supplies (parts) that are used up or consumed in the process of their use. When consumable spare parts fail, they are removed and disposed of. These types of items are managed with an Economic Order Quantity (EOQ) inventory model. The Air Force attempts to achieve target fill rates at each stockage location on these items. Reparable or recoverable items are those supplies or spare parts that, when they fail, can be removed from the weapon system, repaired, and returned to serviceable condition to be reused on the same or a like weapon system. When compared to EOQ items, the mean usage, or demand rate, of reparable items is low, while variability of demand and cost of the item are relatively high. There are many different inventory management models which are used or being developed for use with reparable items.

The Air Force is extremely interested in insuring that it obtains the ability to wage war with stock levels that will have a high probability of supporting this objective without being wasteful of dollar resources. It therefore requires and uses recoverable inventory models that will compute stock requirements to meet certain performance goals of the mission. The Air Force Logistics Command (AFLC) currently uses the METRIC inventory model in its D041 (AFLC Recoverable Items Requirements Computation System) and D028 (AFLC Central Leveling System) to minimize backorders across all bases and the depot for AFLC managed reparable stock items (20). The Wartime Assessment and Requirements System (WARS), currently under development at AFLC, will optimize stock level requirements by weapon system at all bases and the depot to achieve a target performance level at each base (e.g., at least 80 percent of aircraft available at all times) for a given scenario while minimizing the cost of stock to be purchased. WARS will also maximize performance given that there is a limited amount of funds to spend (17:1-4). Models such as Dyna-METRIC (a RAND model) (22:1-3) and WARS can also be used to predict what levels of performance can be expected for a squadron of aircraft given a likely wartime (or peacetime) scenario and existing stock level quantities. Inventory models then become very important tools to achieve desired performance while minimizing stockage costs (stock compu-

tation mode) or for prediction of performance, given levels of stock already on hand (performance mode).

This research deals with the contemporary recoverable item inventory management models and their use in the Air Force inventory management. More specifically, the research examines the sensitivity of certain model outputs to changes in a specific input parameter. The input parameter is the variance-to-mean ratio (V/M), which is the measure of dispersion (or uncertainty) of demand about the average or mean demand. It is simply the ratio of the variance of an item's demand distribution to its mean. The model outputs are dollars of stock necessary to achieve desired performance goals (the probability of having at least the desired number of available aircraft), and scenario performance (probability of having at least the desired number of aircraft available), given a certain stockage policy (i.e., quantities and mix of spare parts).

Recoverable item models cover the spectrum of complexity starting with the early conventional pipeline models all the way through the WARS model currently being developed by AFLC. In this spectrum of models used or being studied for use by the Air Force are the Base Stockage, METRIC, MOD-METRIC, LMI Availability, and Dyna-METRIC models. These models will be dealt with and explained in some detail in the literature review chapter of this report. They are mentioned here to point out that all these models assume

that stock usage demand distributions are stochastic, that is, they have a mean and a distribution about a mean. All these models require a variance-to-mean ratio to express this demand variability.

The conventional pipeline model as used in the Standard Base Supply System (SBSS) utilizes an implied variance-to-mean ratio of three.

The remaining models require a variance-to-mean parameter as an input to the model. The distributions used are generally of two types: Poisson or compound Poisson. By its very nature the variance and mean of a simple Poisson are equal and so the variance-to-mean ratio equals one. The Poisson distribution is a logical distribution to use for demands per unit of time (λ) since the inverse of λ, ($\frac{1}{\lambda}$), which can be considered the mean time between failures, generally exhibits properties of the exponential distribution. According to Sherbrooke (24:8-9), it is also realistic to assume that recoverable item demand distribution will exhibit a compound (specifically logarithmic) Poisson process. The logarithmic Poisson is obtained by considering batches of demand where the number of batches is Poisson and the number of demands per batch is logarithmic. The logarithmic Poisson is almost identical to the negative binomial for variance-to-mean ratios of one to three. Since the negative binomial is convenient to compute it is used in most model algorithms.

Problem Statement

There is much concern over how the variance-to-mean ratio is determined and what value should be used for any given item of stock (27). More basically, however, there is a need to understand just how sensitive the outputs of the recoverable item models are to the variance-to-mean ratio. The outputs of basic concern are: 1) dollars of stock necessary to support a given level of performance (e.g., 80 percent confidence level that there will be no more than 10 percent of the aircraft Not Mission Capable Supply [NMCS]), and 2) performance levels (e.g., confidence level of achieving no more than 10 percent NMCS rate) given a certain level of input stock.

Research Objectives

The objective of this research is to give the Air Force logisticians some insight into just how their decisions in establishing the variance-to-mean ratio for each recoverable item affect the cost and performance outputs of the contemporary analytical recoverable item models they are working with. The criticality of their decisions in this regard depends upon the sensitivity of outputs to the variance-to-mean ratio.

Research Questions

(1) When computing stock levels to achieve given

performance parameters, using contemporary recoverable item models, is the model output of required dollars of stock sensitive to changes in the variance-to-mean ratio, all other things remaining the same?

(2) When computing the performance (e.g., probability of achieving a desired NMCS level) of a group of aircraft, given a likely scenario and a fixed quantity and mix of spare parts, is the model output of mission performance sensitive to changes in the variance-to-mean ratio?

Scope

A state-of-the-art recoverable item inventory model and a validated data base formed the foundation of this research. Dyna-METRIC, a RAND Corporation creation, was the analytical inventory model selected. It was chosen because of its versatility. The model will handle multi-indentured spare parts (line-replaceable units [LRU]/shop replaceable units [SRU]) in a multi-echelon stockage and repair mode. The echelons which can be modeled are bases, centralized intermediate repair facilities (CIRF), and depots. Dyna-METRIC will function in both the stock and performance computation modes. However, because the model will work in the stock computation mode only if a full SRU cannibalization assumption is made, the scope of the research is confined to model usage in the full cannibalization mode. Cannibalization is the act of removing a

spare part from one unit to insert into another unit to make the latter operable. Dyna-METRIC also allows the researcher to change flying hour programs, sortie rates, and aircraft levels at each base modeled, both in peace and wartime scenarios, making it a very powerful tool in modeling realism into scenarios.

The data base is a fully validated F-16 data base provided by Ogden Air Logistics Center (OOALC). Spare parts demand data and the other scenario specifications contained are based on either actual experience or forecasts based on actual experience (11). It is a relatively static peacetime scenario with minimal changes in flying program parameters. Repair, delay, and transportation times are unchanged for each individual spare part over the entire scenario. All parameters except the variance-to-mean ratio (in stockage and performance computations) and the stock levels (in performance computations) will remain as supplied by OOALC.

Interpretation of results will be generalizable to the F-16 weapon system. However, results will have implications for managers of spare parts for any weapon system. This research should provide the Air Force with an indication of whether or not it should be concerned with how it establishes and uses variance-to-mean ratios. This will depend on the sensitivity of model outputs (e.g., stockage costs and aircraft availability) to changes in the variance-to-mean ratio.

CHAPTER II

LITERATURE REVIEW

Introduction

This chapter will deal in detail with the various models that have been designed to make the Air Force inventory system more effective and, at the same time, more efficient. As discussed previously, the inventory system is made up of two classes of items: consumables and reparables. Presently no inventory model exists which handles both consumables and reparable items jointly. The consumable items are managed under the EOQ policy and will not be discussed here. This study will concern itself with reparable or recoverable items since it is an area with many unanswered questions and because recoverables represent a large portion of the Air Force inventory budget. In 1982 the investment in spare parts will be approaching \$2.5 billion, approximately 95 percent of which will be recoverable items (4:4).

This chapter will begin with a discussion of current theory and then discuss performance measures used to evaluate a supply system. Following the performance measures will be a discussion of the various models developed to aid in management of the supply inventory and lastly, a discussion of previous variance-to-mean ratio studies and

current Air Force variance-to-mean ratio estimates.

Inventory Theory

(S-1,S) Inventory Policy. The Air Force uses a (S-1,S) policy to reorder recoverable parts. Any inventory system is designed to answer two questions: when to reorder and how much to reorder. The (S-1,S) policy has been found to be the most cost effective (8:1). The decision for the recoverable system is at what level to set S, the spare stock, for protection against stockouts. Any time the base stockage level drops below S an order is placed at the next higher echelon to bring the total of stock on-hand, plus on-order, minus backorders, equal to S (7:1).

The supply process for reparable items is complicated, with many participants. Figure 1 represents a typical supply system for recoverable items. Essentially what happens is that when an item fails, the base determines whether or not it can fix the item. If the item can be repaired at base level it traverses through the base repair cycle and is returned to serviceable stock. If the item is not reparable at base level, it is forwarded to the next echelon of repair, while concurrently a requisition is placed on the next echelon to bring the base stock level back to the S level.

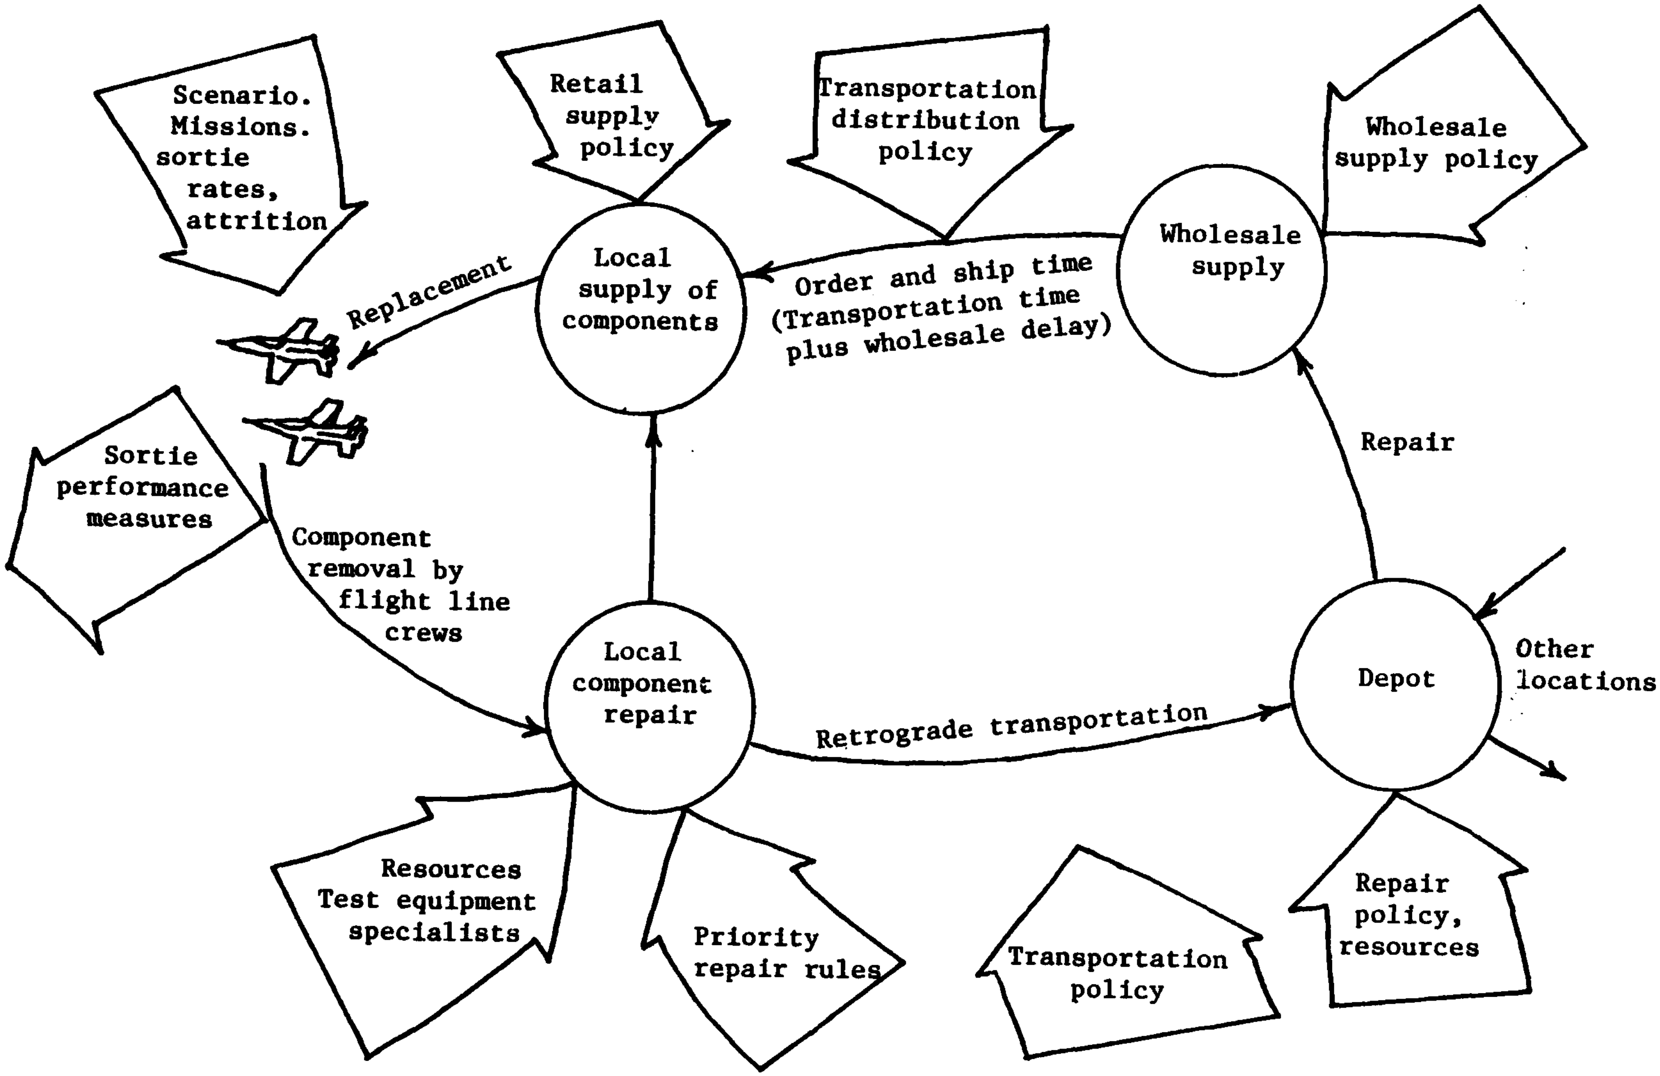

Figure 1. Aircraft Component Support Network (22:13)

Palm's Theorem (Steady-State). Feeney and Sherbrooke (7:2-4) were able to take a queueing model developed by C. S. Palm and apply it to the Air Force inventory system. They demonstrated that when demand is Poisson, "the number of units in resupply in the steady-state, x, is also Poisson for any distribution of resupply." The Poisson state probability depends not on the resupply distribution but only on the mean of the resupply distribution for any member of the Poisson family. A theorem developed by Feeney and Sherbrooke (7:3) states that when demands are Poisson with a mean rate of λ and the resupply time is an arbitrary probability distribution with mean T, the probabilities of x units in resupply are given by the Poisson distribution with parameter λT, i.e.,

$$p(x|\lambda T) = \frac{(\lambda T)^x e^{-\lambda T}}{x!}. \qquad (1)$$

Assumption of Models. In order to be able to predict the behavior of any item the following assumptions were made by Brooks, Gillen, and Lu (1:6-7):

(1) Whenever an operable asset is required, the customer must turn in the old carcass to base supply. The customer will then receive either an asset from base supply's on-hand stock or a due out from base supply. The customer will eventually get an asset for this due out

that is either a repaired item that went through the base repair system or an item sent from the depot. In all cases a one-for-one exchange is made with the next higher echelon of supply. The lower echelon returns the defective part and receives either an operable asset or the promise of an operable asset.

(2) If a replacement item is not available immediately from base supply, the item is backordered and the customer receives a due out from base supply.

(3) The number of demands during a time interval is a Poisson random variable, statistically independent of the number of demands that occur in any other non-overlapping period.

(4) The demand is stationary over an interval of time. The probability distribution of the random Poisson variable depends only on the length of the time interval. Identical time periods have identical probability distributions. There is no trend or cycling of demand.

(5) It is assumed that the length of time required to repair an item is independent of the number of demands.

Inventory System Performance Measures

Fill Rate. The fill rate is one performance measure used by the Air Force to evaluate the supply system. The fill rate is determined by taking the ratio of the number of items issued over the total number of demands placed on

base supply for a certain, fixed time period (1:2). Using Palm's theorem it is possible to predict the fill rate for a given stock level. According to Brooks, Gillen, and Lu (1:10), the fill rate for item i is the proportion of demands "that can be filled from on-hand assets and is equal to the proportion of time that on-hand assets are positive." Mathematically, the fill rate is given by:

$$FR_i = \sum_{d=0}^{S_i - 1} p(d|\lambda_i T_i) \qquad (2)$$

where FR_i = fill rate for item i,

S_i = stock level for item i, and

d = number of items in resupply.

Ready Rate. The ready rate is the probability that an item has no backorders when observed in a random point in time. It is the inclusive probabilities of the number of items in resupply (d) being no greater than S, the stock level and is given by:

$$RR_i = \sum_{d=0}^{S_i} p(d|\lambda_i T_i) \qquad (3)$$

where RR_i is the ready rate for item i (8:4-5).

Expected Backorders. A backorder occurs whenever the number of assets in resupply (d) for an item exceeds the quantity of available spare assets (S). Brooks, Gillen, and Lu (1:11-12) define the expected backorder performance measure for item i, $E(B_i)$, as:

$$E(B_i) = \sum_{d=S_i+1}^{\infty} (d-S)\,p(d\,|\lambda_i \tau_i). \qquad (4)$$

Non-Mission Capable Supply (NMCS). Another performance measure used is NMCS. NMCS is a measure of the number of aircraft that are not mission capable or not operational because of a lack of spares. The NMCS measure can be predicted under two different assumptions: full cannibalization (full cann) or no cannibalization (no cann). Cannibalization is the act of replacing a defective part with an operable part taken from another aircraft. Both the full cann and the no cann assumptions are extreme, the true cannibalization rate probably lies between the two extremes. In reality, a unit will probably cannibalize from one airplane to another but the unit will not cannibalize all the broken items into the fewest number of aircraft. Full cannibalization is not possible because some items suffer a high damage rate when cannibalized and/or the manpower requirements for full cann are prohibitive.

a. No Cann NMCS. A model developed by the Logistic Management Institute (3:44-45) gives a measure of aircraft availability. The complement of availability is the NMCS rate. The probability of any aircraft being available, at a random point in time, can be expressed as:

$$p(\text{available}) = \prod_{i=1}^{n} \left[1 - \frac{E(B_i)}{F \cdot QPA_i}\right]^{QPA_i} \tag{5}$$

where $E(B_i)$ = expected backorders for item i,

n = total number of items,

F = total number of airplanes, and

QPA_i = quantity per aircraft for item i.

The number of available aircraft (AA) is given by:

$$AA = p(\text{available}) \cdot F. \tag{6}$$

The expected number of aircraft NMCS, E(NMCS), or not available, is given then, by:

$$E(NMCS) = F - AA. \tag{7}$$

b. Full Cann NMCS. The operational rate (OR) is the probability that, at any point in time, there will be no backorders. Then, OR is the probability that no air-

craft lacks an essential part, i.e., there will be no NMCS aircraft. The operational rate can be expressed as the product of the ready rates:

$$OR = \prod_{i=1}^{n} RR_i = \prod_{i=1}^{n} \sum_{d=0}^{S_i} p(d|\lambda_i \tau_i) \qquad (8)$$

where n = number of items,

d = number of items in resupply, and

S = stock level (1:4,12).

The previously discussed operational rate assumes zero or no cannibalization. It is possible to express the operational rate as a function of the number of airplanes used as source of supply (K). For example, an operational rate with K = 3 is the same as saying that the parts of three airplanes can be considered as belonging to the spare parts inventory. Mathematically, the operational rate with cannibalization can be expressed as:

$$OR \text{ (with cann)} = \prod_{i=1}^{n} \sum_{d=0}^{S_i + (QPA_i \cdot K)} p(d|\lambda_i \tau_i) \qquad (9)$$

where K = number of aircraft cannibalized (1:4,13).

In order to determine the number of NMCS aircraft, it is assumed that in full cannibalization broken parts are

consolidated onto as few aircraft as possible. The probability of K or less NMCS aircraft is the same as equation (9), the operational rate with K aircraft available for cannibalization (1:13-14). Thus:

$$p(\text{K or fewer NMCS aircraft}) = \prod_{i=1}^{n} \sum_{d=0}^{S_i+(QPA_i \cdot K)} p(d|\lambda_i T_i). \qquad (10)$$

The probability of NMCS is then used to determine the expected number of NMCS aircraft in the following equation:

$$\text{Expected NMCS} = \sum_{K=0}^{\infty} \left[1 - \left(\prod_{i=1}^{n} \sum_{d=0}^{S_i+(QPA_i \cdot K)} p(d|\lambda_i T_i) \right) \right]. \qquad (11)$$

The Expected NMCS equation can be generalized into:

$$E(NMCS) = \sum_{i=0}^{F} Kp(K) \qquad (12)$$

where p(K) is the exact probability that K aircraft are not available. P(K) is the operational rate with K cannibalizations minus the operational rate with K-1 cannibalizations. F is the total number of aircraft (1:14-15).

Palm's Theorem (Nonsteady-State)

Palm's theorem is used widely to describe the arrival process. The main reason for its widespread use is because the Poisson distribution, used in Palm's theorem, is a close approximation to the real world arrival processes (2:1).

Although the real world arrival process might approximate a Poisson distribution, it is rare to find the arrival process to remain in steady-state (2:1). Muckstadt (18:5-7) found that a steady-state assumption used in a wartime dynamic environment would lead to inadequate supply support when it was needed most.

In 1976, RAND developed a technique that allowed the classical form of Palm's theorem to be generalized to a dynamic or nonsteady-state arrival process (2:2). The number of items in steady-state resupply (λT) is given by:

$$\lambda T = \bar{d} \cdot T \qquad (13)$$

where $\bar{d}$ = average demand, and

T = average resupply time for the steady-state model of Palm's theorem (13:7).

The number of items in resupply, at time t, for the nonsteady-state Palm's theorem is given by:

$$\lambda T(t) = \int_0^t d(s)\overline{F}(t,s)ds \tag{14}$$

where $\lambda T(t)$ = items in resupply at time t,

$d(s)$ = daily failure rate at time s,

$\overline{F}(t,s)$ = probability that a component that arrived at time s, is not out of repair by time t, and

ds = interval of time centered at s (13:11-12).

Once the value for λT is determined, the probability of the number of units in resupply at time t, is given by Palm's theorem (13:12):

$$p(x|\lambda T(t)) = \frac{\lambda T(t)^x e^{-\lambda T(t)}}{x!} \tag{15}$$

The difference between the steady-state and the dynamic models is their sensitivity to the distribution of the repair time. Whereas both models need to know the mean of the repair time, only the dynamic model needs to know the distribution of repair time. If the distribution is not known and is approximated by a degenerate repair time distribution (i.e., repair time is constant), the model gives conservative results (2:16,27). The use of a steady-state model in a dynamic environment will not capture the dynamic demands and will misallocate resources. But a

nonsteady-state model with an assumed constant repair time will result in more items in repair during any surge in demand, but fewer in repair following the surge than if a distribution around the mean (e.g., exponential) is known (2:14).

Models

Conventional Air Force Base Level Stockage Model. This model, like all others, computes S, a pipeline quantity. A rule of pipelines is that if S is not large enough to fill the pipeline, then the pipeline will fill itself by taking parts from the aircraft. The pipeline must be full (13:3).

Presently, this is the model used by most bases in the Air Force to determine their requirements, although it is being phased out and replaced with a version of METRIC. The weaknesses of this model are that it treats each item independently of all other items, does not take cost into account, nor does it give any indication of weapon system performance. S, the quantity demanded for the pipeline is given by:

$$S = DDR \times (PBR \cdot RCT) + ((1 - PBR) \cdot OST) \qquad (16)$$

where DDR is the average number demanded per day for that item. PBR is the percentage of base repair for that item.

PBR will be a percent between 0 and 100 and represent the percentage that were repaired on base. 1-PBR is the complement of PBR, the percent that were sent off-base to be repaired. The 1-PBR rate is also called the Not Reparable This Station (NRTS) rate. RCT is the average repair cycle time for that item. Currently, the Air Force uses either a five period moving average or a standard value of five days for RCT, whichever is smaller. OST is the order and ship time, the time required to receive the item from a higher echelon of supply. The OST value is currently set at a historical average value for each base and is usually on the order of 20 days (29).

Because the Conventional model algorithm assumes that the daily demand rate is constant, it does not provide for demand variability. To compensate for this problem the Conventional model adds to S a safety level of $\sqrt{3S}$. The model assumes a normal distribution for the quantity in resupply and sets $\sqrt{3S}$ equal to one standard deviation, σ. By assuming a normal distribution and $\sigma = \sqrt{3S}$, the model attempts to achieve an 84 percent fill rate. That is, 84 percent of the demands will be filled from on-hand serviceable stock. By assuming $\sigma = \sqrt{3S}$, the model assumes a variance-to-mean $\frac{\sigma^2}{S}$ equal to 3 for all items (29).

<u>Base Stockage Model</u>. As compared to the Conventional model, which looks at each item separately, the RAND Cor-

poration developed, in the mid-sixties, a method of looking at all items at an individual stockage location (e.g., a base). The Base Stockage Model was a first attempt to look at all the items in base supply and make trade-offs among the items to maximize a system objective subject to a cost constraint (8:10-11).

The system objective of the Base Stockage model is to minimize the number of backorders E(B) subject to a dollar constraint. A backorder occurs whenever the number in resupply exceeds the stock level (S) of that item. The expected number of backorders is given by equation (4). A decrease in backorders is produced by increasing the inventory levels from S to S+1. The amount of backorder reduction, as formulated by Faucett and Gilbert (6:25-26), is given by:

$$E(B_S)-E(B_{S+1}) = \sum_{x=S+1}^{\infty} p(x|\lambda T). \qquad (17)$$

The model performs a marginal analysis for all the items at that base and determines what mix of stock will give the lowest backorder level with a given quantity of money. It assumes a zero stock level for every item at the base and then asks what item would give the most fill protection per dollar of stock (9:14-15).

The data that was used by Feeney and Sherbrooke was found to have a large variability in demand (9:19). To account for this variability they devised a parameter they call a variance-to-mean ratio. This term (8:8) is "estimated from two or more periods of cross sectional data." The value for variance-to-mean ratio is inserted into the algorithm to achieve the correct probability distribution associated with that variance-to-mean (8:8).

The variance-to-mean ratio is a function of the Poisson distribution. A simple Poisson has a variance-to-mean ratio of one. A simple Poisson can be visualized as a single customer arrival in each time interval. That customer places only one demand and the time distribution is exponential. A compound Poisson will have a variance-to-mean ratio greater than one. The compound Poisson can be visualized exactly like the simple Poisson except that a customer can place multiple demands on the system according to an independent discrete distribution. The compound Poisson distributions are the most general of the memoryless distributions. In other words, the number of demands that occur in one time period does not influence the probabilities in any other time period. The compound Poisson seems to be a better description of demands for the Air Force recoverable spares than does the simple Poisson. Customers do occasionally place more than one demand at any one time. Some possible reasons for multiple demands

are: 1) sympathetic replacement or undetected malfunction, 2) initial wearout, and 3) damage during installation (7: 4-6).

The Base Stockage model faces the same dilemma that the Conventional model faces, how to deal with demand variability. Whereas the Conventional model made an arbitrary decision, the Base Stockage model has incorporated the variance-to-mean ratio into its algorithm.

In summary, the Base Stockage model, by incorporating the systems approach to inventory control, allows the supply manager to select an alternative that is the most cost effective for that investment level. He or she will immediately know the performance level of the system. Because this model optimizes, a supply manager could attain, theoretically, the same performance level with one-half the investment (8:23).

METRIC. Although the Base Stockage model was an improvement over previous methods of managing a base inventory, it was never implemented in the Air Force. The Base Stockage model ignored the depot or any other base in its computations. It optimized only at one base at one time. After testing, but before implementation, Sherbrooke was able to expand the Base Stockage model to include higher levels of supply. This new model is referred to as Multi-Echelon Technique for Recoverable Item Control (METRIC).

METRIC uses essentially the same logic as the Base Stockage model but expands it to include the complex base-depot relationship. By optimizing over the entire supply system, the METRIC model avoided many of the pitfalls of the Base Stockage model in a more cost effective manner. The model is able to both compute requirements and to redistribute stock more appropriately (25:122). The objective of METRIC is again to minimize expected backorders over the specified items subject to the investment constraint. Depot backorders are a factor only as they affect base backorders (19:473).

The Air Force is in the process of implementing METRIC more fully into its inventory management system. Presently the D041A uses METRIC to compute requirements. The Conventional model is being phased out at the base level to be replaced with METRIC logic, and the D028A is also partially converted to METRIC logic (20). The benefits of METRIC are: 1) its ability to compute stock levels across all items, and to distribute those items in a manner that optimizes system performance within budgetary constraints, 2) because METRIC is so similar mathematically to the Base Stockage model all the experience gained from the Base Stockage model is directly relevant to METRIC, 3) the model uses past data in combination with future requirements to accomodate buildups or phase outs, 4) by using a "general demand prediction procedure", METRIC allows a

smooth transition from initial support to follow-on provisioning, 5) the model allows resupply or repair parameters to be varied for sensitivity analysis of system performance, and 6) the model allows management to pursue different support policies for different weapon systems (25:124).

As was the case with the Base Stockage model, Sherbrooke assumes that demand for aircraft spares follows more closely a compound Poisson. This time the compound Poisson is the logarithmic Poisson. The logarithmic Poisson retains the tractable properties of the Poisson process, but allows greater flexibility because of the additional parameters. The logarithmic Poisson can be visualized as a process where a batch of demands arrives according to a Poisson distribution and the number of demands per batch follows a logarithmic distribution. Sherbrooke found the state probabilities (probabilities of n demands during time t), for the logarithmic Poisson to be negative binomial. The negative binomial is a particularly easy distribution to work. For variance-to-mean ratios greater than one and less than three, the state probabilities for the two distributions are essentially identical. For example the probability of x demands in the specified time period is given by the following negative binomial distribution:

$$p(x|\lambda) = \left[(k+x-1)!/(k-1)!x!\right] \cdot \left[(q-1)^x/q^{k+x}\right] \quad (x = 0,1,\ldots, \; q>1, \; k>0) \qquad (18)$$

where $\lambda = k(\ln q)$,

λ = mean demands per time period, and

q = variance-to-mean ratio (24:8-9).

An AFIT study done recently on the failure of Inertial Measurement Units found the negative binomial distribution to best describe the failure process (16:35-36). As was the case with the Base Stockage model, mean demand is assumed to be stationary for the computation period of concern (25:129).

In summary, METRIC is an extension of the mathematical model developed for the base stockage inventory system. It is an improvement over the Base Stockage model because it describes both the bases and the depot and makes trade-offs across items and across locations to minimize backorders across the whole system.

MOD-METRIC. An extension of METRIC was advanced by Muckstadt. He noticed that METRIC tended to concentrate more heavily on inexpensive sub-components because it was able to decrease the backorder level more by buying these items. MOD-METRIC overcomes this problem by establishing an indenture relationship between components and sub-components. The component is called a line replaceable unit (LRU), and is the end item that failed. An LRU is removed at the flight line and taken back to the shop where it is repaired by removing the defective sub-component and re-

placing it with a new one. The sub-component is called a shop replaceable unit (SRU). MOD-METRIC assumes that LRUs will ground aircraft, while an SRU backorder merely delays an LRU from being repaired. The algorithm is designed to minimize backorders on LRUs soley subject to an investment constraint (20:472-475).

The difference between METRIC and MOD-METRIC models is the manner in which the average resupply time ($\mathcal{T}_i$) is computed. METRIC expresses $\mathcal{T}_i$ as:

$$\mathcal{T}_i = r_i B_i + (1-r_i)(A_i + \delta(S_o)D) \qquad (19)$$

while MOD-METRIC expresses $\mathcal{T}_i$ as:

$$\mathcal{T}_i = r_i(R_i + \Delta_i) + (1-r_i)(A_i + \delta(S_o)D) \qquad (20)$$

where r_i = the probability that a failure isolated to an SRU will be repaired at base,

B_i = average base repair time,

A_i = average order and ship time,

$\delta(S_o)D$ = expected backorders/expected daily demand at the depot,

R_i = average repair time at base if SRU_i is present, and,

Δ_i = average delay in base repair due to the unavailability of a needed SRU.

MOD-METRIC is concerned solely with the minimization of LRU backorders subject to a dollar constraint. By dividing the average repair time (B_i) into repair time with the SRU (R_i) and the average delay in repair caused by an unavailable SRU (Δ_i), MOD-METRIC minimizes the LRU backorders and provides an optimal stock level for both the LRU and the SRU (20:476).

LMI Availability Model. Although previous inventory models have provided a measure of the supply system performance, there was the desire by many (e.g., Congress and the DOD) to have a model that expressed the military capability of the weapons system (22:1). The LMI Availability model, a no cannibalization model, was developed to maximize aircraft availability by the algorithm given in equation (5).

The LMI Availability model, a further extension of METRIC (3:8), "is simply inserted into the METRIC program." While the METRIC system minimizes expected backorders, the LMI Availability model minimizes the expected number of NMCS aircraft and consequently maximizes availability (3:54). METRIC has as an objective function the minimization of a sum of its terms while the LMI Availability model minimizes the product of its terms. "This is quite a difference in objectives [3:58]."

Dyna-METRIC. Although models exist that predict the amount of backorders we can expect or the amount of aircraft we can expect to have available, concern was expressed that all those models might not be able to answer the paramount question: How well does that weapon system perform in combat? All the previous models assumed that the weapon system was working in a constant environment. These models assume that all input variables will remain in a "steady-state" condition. There was no way of predicting in a combat environment that contained much turbulence whether logistics could still support the weapon system. The steady-state models contain no method for determining, for example, the effect a disruption in the transportation line would have on weapon support, nor what would be the effect of limited maintenance capability. These models had no ability to adjust to the complex, dynamic situations in which modern weapons systems could become involved (10: 12-13).

During peacetime, any shortfall can be corrected over time, with no loss of life or territory. But in wartime, a mistake is not as easily forgiven. America enjoyed the luxury, in World War II, of having the time to respond to the threat. Future wars that America might become involved in are perceived to be of short duration of an extremely high intensity. America will not have the time to respond as in World War II. If the logistics support

required in such a war could be resupplied instantaneously, there would be no problem. But today's logistics manager will need time to build the logistics capability. He or she must insure that either there will be no disruption in logistics support or the combat unit is organically capable of supporting its wartime mission. Evidently, it is impossible to guarantee anything and the safest option is to have each unit organically self-sufficient.

The dilemma encountered can only be resolved if today's logistics manager is able to predict how logistic support affects operational capability. He or she must know how the level of maintenance support, transportation support, resupply of parts, or any of the whole myriad of logistics components affect operational performance. If a shortfall exists, but is not detected, the possibility exists that the mission is jeopardized. Because of the lead-time required to position equipment, manpower, and the other logistics components in place, it is imperative that any inventory model be able to predict the correct requirement to fix any shortfall in time to take corrective action (22:1-2).

Dyna-METRIC is a RAND developed inventory model designed to allow the logistics manager the ability to look at any potential combat environment and determine the shortfalls caused by inadequate logistical support. This ability to deal with the transient qualities of

combat (e.g., sortie rates, phased arrival of repair capabilities, etc.) is the strength of Dyna-METRIC over the steady-state models. That coupled with its capability to calculate availability of aircraft, allows for it to be of great use to the logistics manager (13:3,5).

As before, Dyna-METRIC is an extension of previous models. METRIC, MOD-METRIC, and the availability models are all incorporated into Dyna-METRIC with modifications to account for the dynamic environment.

Dyna-METRIC took the work that had been done on METRIC, which assumed steady-state, and "derived similar results for a time varying service and demand process [14:5]." The model predicts how component availability will affect the forces combat capability measured in the number of Fully Mission Capable (FMC) aircraft based on the support system's resources, transportation capability, and repair capability (22:2). Because Dyna-METRIC is an analytical model, it can be operated either to compute stock levels or to determine performance. In both modes the operator specifies the probability he or she desires to be "sure" of having a specified level, or less, of NMCS aircraft. For example, in the work done on this thesis, the authors desired an 80 percent probability that there would be 10 percent or fewer NMCS aircraft at each base. The model would look at each base at the specified time and determine what probability existed of there being 10 percent

or fewer NMCS aircraft. If the model determined that the likelihood of achieving the 10 percent NMCS goal was less than 80 percent, it would recognize it as a problem. In the performance evaluation mode, it would identify the parts that caused the system to not achieve the desired probability. While in the stock computational mode, it would recommend an efficient marginal stock purchase, if necessary, to improve aircraft availability (22:12).

The objective of Dyna-METRIC is to avoid the loss of aircraft missions because of a shortage of functioning components. This can only happen when the number of components in repair and transportation is less than the local spares supply of these components. At any one point in time, one could find some serviceable items in inventory, some items awaiting repair, some items being repaired, some items in transit to and from higher echelons of resupply, and some items partially repaired but awaiting parts before further repair can commence. Each of these states is a pipeline and contains a portion of the total inventory for that item. If any of these pipelines is not "filled" with spares, there can be a corresponding hole in an aircraft which will affect mission accomplishment for mission essential components (13:3). The central computation of Dyna-METRIC is to determine the number of items for each component in each pipeline (22:14). The model then takes the number that it (probabilistically) expects in the total

pipeline and compares that with the available stock at each location. It then forecasts (22:16) "how the nonserviceable components would (probabilistically) generate backorders (or aircraft holes) and how those holes would affect the number of available aircraft," for both the full and no cannibalization cases. Both cannibalization modes have a mean and a standard deviation computed but the aircraft probability distribution represents full cannibalization (22:16).

The algorithms for Dyna-METRIC are similar to other models encountered so far. Expected backorders, at a point in time, EB(t) is given by:

$$EB(t) = \sum_{d=S(t)+1}^{\infty} (d-S(t))p(d|\lambda T(t)) \qquad (21)$$

where d = number of components in resupply at time t,

S(t) = stock level at time t, and

$\lambda T(t)$ = average number in pipeline at time t.

It is obvious that Dyna-METRIC uses the same expected backorder equation (4) as the earlier Base Stockage model, adapted for a nonsteady-state environment (13:28-29).

The number of NMCS aircraft at time t, (EN(t)), for a no cannibalization scenario is given by:

$$EN(t) = F(t)\left[1-\prod_{i=1}^{n}\left(1-\frac{EB_i(t)}{F(t)}\right)\right] \qquad (22)$$

where $F(t)$ = number of aircraft at time t,

n = number of items, and

$EB_i(t)$ = number of backorders for item i at time t (13:33-34).

Again, the equation is exactly like equation (7), but adapted for a nonsteady-state environment.

A full cannibalization scenario (13:35-36) assumes that shortages are consolidated to make the smallest expected number of non-mission capable aircraft. The expected number of non-mission capable aircraft, $EN(t)$, is given by:

$$EN(t) = \sum_{j=0}^{F(t)-1} 1-p(j) \qquad (23)$$

where $p(j)$ = the probability that the number of NMCS aircraft is equal to or less than j. Again this is an adaptation of equation (11) for a time varying scenario. The NMC distribution function under the full cannibalization is:

$$pN_j(t) = p(j)-p(j-1). \qquad (24)$$

Although Dyna-METRIC is an attempt to model the real world more accurately, it still has its limitations (22:34-43):

1. The repair process is stationary. Dyna-METRIC assumes that if a repair team is present, adequate repair resources exist to assure the average repair time is relatively constant. There will be no parts in queue awaiting repair.

2. The sortie rate is not constrained by flight line limitations.

3. Component demand rates vary with the flying hours only.

4. Aircraft at any one base must be "semi-homogeneous". The model assumes for full cannibalization that all aircraft, their sorties, and their components are fully interchangeable.

5. Components must be fully tested before repair, or reordering can begin.

6. All levels of resupply are assumed to have identical repair processes. The repair times are identical as are the sub-component demand rates.

Dyna-METRIC has the capability to assign the appropriate variance-to-mean ratio to any LRU. In the future, variance-to-mean ratios will be permitted on the SRU data input (16). One can specify a variance-to-mean ratio from zero to infinity. Assuming that the variance-to-mean

ratio equals one, the probability distribution of x items in resupply is a Poisson random variable, given by equation (1). For variance-to-mean ratios greater than one, Dyna-METRIC uses the same negative binomial distribution as equation (18).

Sensitivity Analysis of the Variance-to-Mean Parameter

Every inventory model requires that consumption demand be expressed in terms of a mean demand and the variability of that demand about its mean. Beginning with the Base Stockage model, it became possible to analyze the supply system by varying input parameters (8:14).

Experimenting with three different variance-to-mean ratios, Feeney and Sherbrooke (9:21) found "stockage requirements to be very sensitive to assumption of demand variability." An optimization based on an incorrect variance-to-mean ratio usually would lead to "substantial degradation in performance [8:21]." They found, for example, that an assumed variance-to-mean ratio of one (simple Poisson) would lead to a fill rate of .886 if the true variance-to-mean ratio is one. But, if the assumed variance-to-mean ratio is one and the true variance-to-mean ratio is two, one would have a fill rate of .845. In order to bring the fill rate up to .886, a 25 percent increase in stock would be required. The situation is even more bleak if the assumed variance-to-mean ratio is one, while the true

variance-to-mean ratio is three. In that case, a 50 percent increase in stock would be required to achieve the same fill rate (8:21-22). Feeney and Sherbrooke (8:22) conclude that the correct variance-to-mean ratio is "essential for an appropriate allocation of funds." Further, they insist that the variance of all items must be reduced to the minimum possible to reduce stockage requirements and consequently, investment.

The METRIC model assumes that the variance-to-mean ratio is the same for all bases. This assumption allows the demand distribution at the depot to be logarithmic Poisson with the same variance-to-mean ratio, a feature which allows a computational advantage (25:132). Sherbrooke (24:47) investigated the sensitivity of the variance-to-mean ratio in METRIC. With a true variance-to-mean of two and an assumed variance-to-mean of one, he found a 29 percent increase in expected backorders. Inverting the relationship by assuming a variance-to-mean of two, when the true variance-to-mean is one, yields an 82 percent increase in backorders. Because METRIC is cost constrained, each alteration of the variance-to-mean ratio would distribute the assets differently. Evidently, the distribution that is optimal for an assumed variance-to-mean ratio of two is not optimal for a real variance-to-mean ratio of one. Sherbrooke also found that the variance-to-mean ratio became more sensitive as the investment increased.

A complementary study by Sprung (26:8) at AFLC found that as the variance-to-mean ratio increased, the number of backorders increased if the algorithm was constrained by a dollar amount. Estey (5:16), while working on a non-METRIC inventory system, found the cost of the inventory to increase as the variance-to-mean ratio increased.

Current Practices

The questions of what effect various variance-to-mean ratios have on the output of inventory models led the authors to determine what variance-to-mean ratios were currently being used in the Air Force and how they were determined.

The authors interviewed Mr. Victor Presutti, Mr. John Hill, and Mr. Michael Niklas at AFLC/XRS, during the course of this thesis and found them an extremely valuable source of information. In conversation with XRS, the authors found that currently used variance-to-mean ratios vary widely depending on the inventory model used. Besides the inconsistency, there appeared to be only minimal attempts to link the variance-to-mean ratios with the true variability of demand. The variance-to-mean ratio is determined by arbitrary methods (12; 20; 21).

Prior to a 1973 study by AFLC, a standard value of 1.01 was used for the variance-to-mean ratio. It was felt

that 1.01 would not seriously impair the effectiveness of any marginal analysis computations. But more importantly, there was a fear that another value for the variance-to-mean ratio might cause more serious misallocation of resources (28:8). In that study, the historical demand data of items was regressed to find a method of determining the variance-to-mean ratio for each item. A power formula equation of the form

$$\text{Variance-to-mean ratio} = am^b$$

where a and b are regression coefficients, and m is the mean demand, was found to best fit the data. The general relationship was that as the mean increases, the variance-to-mean ratio also increases, but at a decreasing rate (28:17). Currently, the power function is used on the AFLC Recoverable Items Requirements Computation System (D041A), and the AFLC Central Leveling System (D028A)(20).

Far more common than a mathematical technique like the power function is the case where the variance-to-mean ratio is set as a standard figure for the entire system. Currently, the Wartime Requirement Spares Kit (WRSK) and the Base Level Self-Sufficiency Spares (BLSS) set the variance-to-mean ratio at 1.0. Both of these inventory models are methods for establishing wartime requirements (21). A more recent model that uses 1.0 is the Combat Support Capability

Management System (CSCMS) (11). The CSCMS uses Dyna-METRIC as its inventory modelling sub-system and was developed jointly by the Pacific Air Forces (PACAF), AFLC, and RAND (23:iv). According to Pyles (22:28):

> Data and models indicating non-Poisson component pipeline distribution are rarely available. Therefore, most analyses assume the Poisson distribution will be obtained for most components by setting their variance-to-mean ratios to 1.00.

The Logistic Management Institute (LMI) developed the availability model mentioned earlier. This model is used at the Pentagon to forecast aircraft availability for future budget years. Values for the variance-to-mean ratio are set according to the number of years forward the model is used to predict. Close-in years use smaller variance-to-mean ratios, while out years use higher variance-to-mean ratios. The variance-to-mean ratio can range from a value of 1.00 to 5.00 (21).

The most exotic technique for determining a variance-to-mean ratio was developed by AFLC/XRS. In this method, the variance-to-mean ratio is set equal to the square root of the number of bases. The variance-to-mean ratio cannot exceed 5.00. Because of the frustrations XRS feels at not being able to determine the true variance-to-mean ratio, and the variety of variance-to-mean ratios in use, they feel that this technique is as accurate as any. XRS feels that the variance-to-mean ratio is an "extremely sensitive" parameter. But no one knows how sensitive the variance-

to-mean ratio is, or how to compute the true variance-to-mean ratio (21).

In order to reduce the ambiguity of estimating the variance-to-mean ratio, AFLC recently let a contract to determine the best technique to correctly estimate the variance-to-mean ratio (21).

Summary

It is evident that much money ($2.5 billion) is involved in recoverable assets and it is important to the Air Force that this money is managed most efficiently. It was found that an (S-1,S) inventory system is the most cost effective. With the adoption of Palm's theorem to inventory models, it became possible to predict inventory system performance measures (e.g., fill rate, ready rate, expected backorders, no cann availability and full cann availability). It was recognized that the conventional Palm's theorem covered a steady-state environment and a method was needed to address the dynamic environment in which the Air Force might find itself. Palm's theorem was modified to account for this dynamic environment.

Various models were developed to best decide how to spend the money provided to buy spares. The models started simple and became more complex by building on their predecessor. The Conventional stockage model is a fairly simple, non-probabilistic, model that added an arbitrary value for

safety stock. The Base Stockage model is a probabilistic model that uses Palm's theorem and optimizes over the entire system (Base Supply) to minimize backorders. METRIC takes the Base Stockage model one step further and incorporates the depot. MOD-METRIC found that METRIC would buy too many low cost items and set-up indentured relationships so that optimization would only occur on the most critical parts. By making an addition to METRIC, the LMI Availability model captures, as a performance measure, military capability versus backorders. To this point, the environment has been assumed to be steady-state. Dyna-METRIC breaks with this train of thought and looks at the environment as dynamic and varying with time. Dyna-METRIC includes all the previous models and their system performance measures, but in a time-varying scenario.

Previous studies on the variance-to-mean ratio were limited to steady-state models. Using expected backorders as the measure, the variance-to-mean ratio was found to be an extremely sensitive input parameter. It was surprising, then, to find such variability in techniques used to estimate the variance-to-mean ratio.

Conclusion

It is evident that the variance-to-mean ratio is an important parameter. It is included in all inventory models and the analyses done so far have concluded that inventory

models are sensitive to flucuations in the variance-to-mean ratio. One thing that is not known is how sensitive the output of a nonsteady-state model, such as Dyna-METRIC, is to the variance-to-mean ratio.

CHAPTER III

RESEARCH METHODOLOGY

Overview

Three basic ingredients form the foundation of this research. The first is formulation of an experimental design. The design had to be such that the basic research questions could be answered over a relevant range of variance-to-mean ratios. The second necessary ingredient was choice of an evaluation tool. The evaluation tool would have to be a validated "state-of-the-art" recoverable item inventory model capable of handling realistic peacetime/wartime scenarios in both the stock computation and the aircraft performance evaluation modes. The RAND Corporation developed Dyna-METRIC model was selected. Finding or developing a realistic data base and an operational scenario was the third major area. Ogden ALC had developed and refined a PACAF/USAFE F-16 scenario. This highly realistic scenario was selected for this research. The research questions will be answered by executing the experimental design using the Dyna-METRIC model and the Ogden F-16 data base.

Experimental Design

To insure the general applicability of the results of

this research, the research was conducted for a range of variance-to-mean ratios that most Air Force recoverable items exhibit. Since approximately 80 percent (28:20-21; 12) of AFLC managed recoverable items have variance-to-mean ratios between 1.0 and 3.0, this was the area of focus. To insure that trends in experimental results were identifiable, nine variance-to-mean ratios in increments of .25 between 1.0 and 3.0 were used. A variance-to-mean ratio of 5.0 was also used to determine if the general trends of model outputs at 5.0 follow the same pattern as exhibited between 1.0 and 3.0.

The research design can be thought of as completion of a ten by ten matrix as depicted in Figure 2.

V/M Used For Stockage	Variance-to-Mean Ratios Used To Evaluate Performance 1.00 1.25 1.50 1.75....3.00 5.00	Stockage Cost $(millions)
1.00		
1.25		
1.50	Research Question #1	Research Question #2
1.75		
3.00		
5.00		

Figure 2. Matrix Design

The first research question of how changing the variance-to-mean ratios affect item requirement computations (quantities and mix of stock), and therefore cost, is answered by inputting the estimated variance-to-mean ratios observed in the left vertical column into a recoverable item inventory model. The outputs from that model, when run in a stockage computation mode, are quantities, mixes, and costs of spare parts necessary at the bases and depot to achieve a given aircraft availability rate.

While the first research question asks how changing the estimated variance-to-mean ratios affect stockage costs, the second question asks how aircraft performance (availability) is affected if the actual variance-to-mean ratio experienced for an item or group of items is different than the estimated variance-to-mean ratio used to develop stockage. The stock levels computed for each item, both at the bases and depot for each estimated variance-to-mean ratio, are used as inputs to the inventory model. The inventory model is then run in a performance mode, to determine aircraft availability for each actual variance-to-mean ratio between 1.0 and 3.0 and then at 5.0 as shown in the top horizontal row of the matrix. For each set of computed stock levels resulting from the estimated variance-to-mean ratios, there are individual performance computations run for each of the ten actual variance-to-mean ratios. When the ten by ten matrix is complete, there will be 100

individual performance computations.

This design has its strengths and weaknesses. One of the strengths is that all parameters of the realistic scenario are held constant except those parameters of concern: variance-to-mean ratios (in stockage and performance computations) and the stock levels (in performance computations). This enhances the ability of the researcher to discern the cause and effect relationship of changing the estimated variance-to-mean ratios on stockage cost and the effect that erroneously estimated variance-to-mean ratios have on aircraft availability. The matrix design also presents the results in an easily interpretable fashion. Trends should be easily recognized. With the relevant range of variance-to-mean ratios (1.0 to 3.0) broken into .25 increments, the resolution of the output results should be such that trends will be evident. The assumption that all errors in the estimated variance-to-mean ratios are unidirectional is a weakness in this design. While completing the matrix, all the variance-to-mean ratios of all items, are set to the same estimated value in the stock computation mode. Although it is recognized that it is highly unlikely that all errors in the real world will be biased in the same direction at the same magnitude, the results of this research still should be meaningful in gauging the sensitivity of model outputs to changing the variance-to-mean ratios.

Stockage and Evaluation Model

The RAND developed Dyna-METRIC stockage model was used as the evaluation tool for this research. It is a validated "state-of-the-art" recoverable item inventory model with both stockage computation and performance evaluation capability. The problem of formulating data base inputs is greatly simplified when a single model is capable of operating in both modes. The model was made available to AFIT by RAND and was loaded onto the Harris 500 system, located at AFIT.

Dyna-METRIC is a powerful tool that allows the researcher flexibility in setting-up experiments. It is a dynamic model. Aircraft levels and flying hour programs can be changed within any given scenario. This allows realistic changes in both peacetime and wartime scenarios to be modeled, enhancing the credibility of the model's outputs. The only other model with this capability that the Air Force is close to having operational is WARS. But it was not operational at the time of this research and therefore could not be used. Dyna-METRIC is capable of handling indentured relationships among items. If a line replaceable unit (LRU) is composed of several shop replaceable units (SRU), these relationships can be shown. Each LRU and SRU has its own failure (demand) rate. The model is capable of predicting when the LRU will fail and be removed from the aircraft. It is also capable of determining whether a

particular SRU had caused the LRU to fail so that the SRU could be replaced thereby returning the LRU to serviceable condition. In the stock computation mode, the model considers these LRU/SRU relationships and makes decisions to buy whatever quantity and mix of LRU/SRUs are cost optimal to achieve the desired aircraft availability goal. The model is multi-echelon. It will handle up to ten bases and provide full capability to model a Centralized Intermediate Repair Facility (CIRF) or depot. The relationships between bases and CIRF/depot are easily modeled. Shipping, repair, and delay times as well as not reparable this station (NRTS) factors add realism to the scenario. The model is probabilistic. Item consumption (demand) data are input as demand distributions with a mean and variance-to-mean ratio. Model outputs are in terms of probabilities of achieving a desired NMCS rate, number of NMCS aircraft, and sortie generation distributions.

Dyna-METRIC is not without its limitations however. The model can be run in the stock computation mode only under a full LRU/SRU cannibalization assumption. The model automatically assumes that any failed LRU has all of its components SRUs available for use in fixing other failed LRUs. It also assumes that any grounded (NMCS) aircraft has all of its serviceable LRUs available for use to fix other NMCS aircraft. This research, as a result, is confined to the full cannibalization assumption. To gain a

more comprehensive understanding of the effects of changing variance-to-mean ratios, it would have been helpful to experiment at the other extreme of the performance spectrum and assume a no cannibalization policy. RAND currently has plans to develop the capability to do this, but does not know when the enhancement will be available (15). When run in the stock computation mode, the model does not use marginal analysis to decide whether placing stock at the CIRF/depot or any of the bases will do the most good at the least cost in improving aircraft availability. It simply fills the CIRF/depot pipeline and then goes about minimizing the cost of stockage at each base while achieving a desired probability that the NMCS rate will be on target (e.g., p[NMCS rate = 10 percent of aircraft] = 80 percent).

Research Data Base

A USAFE/PACAF F-16 data base (scenario) was obtained from Ogden ALC. This peacetime scenario has been used by Ogden in conjunction with the Dyna-METRIC model for over a year. It is realistic yet relatively simple. It is being used to predict the performance of USAFE/PACAF bases at quarterly increments, given certain stockage levels for each item. Problems in the quantities and mix of spare parts among the bases and depot can be analyzed and corrective action taken. Consumption demand rates, flying hour programs, aircraft levels, repair and delay times, NRTS

rates, and all other data are either based on or are projections from historical data. The data is being updated quarterly, based on the latest information on new incoming stock, buys, or changes to other data parameters such as aircraft flying hour programs and aircraft levels. It is a scaled data base. Since Ogden's objective is to look at quarterly performance increments, they have scaled all time-dependent data so that one day of model time equals 30 days of scenario time. This allows a 36 month look into the future in just 36 days of model time. Running the model is thus much less time consuming with little or no sacrifice in the accuracy of results (11). As a result of this scaling, however, Ogden had to make some input format changes to the Dyna-METRIC FORTRAN coding. So that the Ogden data base could be used in this research, those same changes were incorporated into the version of Dyna-METRIC loaded to the Harris 500.

A comprehensive discussion of the data base is given in Appendix A of this thesis.

Execution of Experiment

Modification to Dyna-METRIC. The Dyna-METRIC model was modified slightly to save time in executing the experiment. FORTRAN statements were added to write to an output file with each stock number and stock quantities for each base and the depot. This file would be written to only when

the model was being run in the stock computation mode. The file could then be manipulated to revise the input stockage portion of the Ogden data base. The data base would then be ready to run in a performance computation mode.

Run Length and Performance Goals. It was decided that both research questions would be answered using the Ogden data base evaluated in a 30 month scenario. The first research question would be answered by running the model and data base in the stock computation mode for 30 months. The objective of the stock computation mode was to achieve an 80 percent probability that there would be no more than 10 percent of the aircraft, at any base, NMCS at any point in time over the entire 30 month scenario. The second research question would be answered by running the model for 30 months using the computed stock levels from question one.

Sequence of Attack. The four corners of the matrix, as outlined in the research design, were completed first. It was felt that these extreme points would indicate whether the range of variance-to-mean ratios considered (1.0 to 5.0) was sufficient to indicate the sensitivity of the performance outputs of the model to the variance-to-mean ratio. The 1.0 versus 1.0 and 5.0 versus 5.0 corners would have to be close to the goal of an 80 percent probability of a 10 percent NMCS rate. The 5.0 versus 1.0 and 1.0 versus 5.0 corners would need to be close to the 100 percent and

0 percent probabilities respectively. If that was not true, modifications to the experimental design would have been undertaken.

Case Study Approach/No Sampling. The nature of experimental execution using an analytical model such as Dyna-METRIC closely parallels a group of case studies. Each of the 100 cells of the ten by ten matrix, when completed, can be thought of as a case study by itself. The mathematical processes are so exacting that for a given set of input parameters only one set of outputs is possible. There is one and only one answer for each cell of the matrix. The answers are, however, given in probabilistic terms with a mean and variability about that mean. This is in contrast to a sampling study using a simulation model. In simulation, a random number generator would produce different sets of input parameters for each run, producing different output results. Several runs would have to be completed for each cell of the matrix and statistical analysis done on those results to reach an answer for each cell. If a truly random sampling was used, there is a very high probability that simulation would come close to but never duplicate the same answer for each cell, no matter how many simulations were run on that cell. The significance of using the analytical model in the multiple case study mode is that it is much less time consuming, allowing the researchers to cover a much broader spectrum of relevant

variance-to-mean ratios in a much more thoroughly detailed experimental design.

CHAPTER IV

RESULTS

Overview

The results of this study will be presented in both tabular and graphic form. Tables 1 and 2 represent the summary of the results in matrix format. Figure 3 addresses the first research question. It graphs the estimated variance-to-mean ratios versus stockage costs. Figures 4 and 5 address the second research question. They show the effects on scenario performance of a stockage policy when the actual variance-to-mean ratio experienced, when the aircraft performed in the field, is not the same as the estimated variance-to-mean ratio used in stock computation.

The analysis of results will focus on interpretation of these tables and graphs. Relationships between variables will be pointed out and analyzed to achieve a full understanding of the results.

Presentation of Research Results

Tables 1 and 2 present the condensed results of the study. The results are displayed in a ten by ten matrix. The far left hand column of each table, gives a value for the estimated variance-to-mean ratio, the value assumed for

V/M USED FOR STOCKAGE	VARIANCE TO MEAN RATIO USED TO EVALUATE PERFORMANCE										STOCKAGE COST $ (MILLION)
	1.00	1.25	1.50	1.75	2.00	2.25	2.50	2.75	3.00	5.00	
1.00	83%	63%	43%	27%	15%	9%	5%	3%	2%	0%	155.15
1.25	93%	83%	67%	51%	36%	24%	16%	10%	7%	0%	172.14
1.50	97%	92%	82%	70%	56%	43%	32%	23%	17%	1%	189.13
1.75	99%	96%	90%	82%	72%	60%	49%	39%	30%	3%	206.33
2.00	99%	98%	95%	89%	82%	73%	63%	54%	45%	8%	223.00
2.25	100%	99%	97%	94%	88%	81%	74%	66%	57%	14%	239.61
2.50	100%	99%	98%	96%	92%	87%	81%	75%	68%	22%	256.00
2.75	100%	100%	99%	97%	95%	91%	87%	81%	75%	30%	272.06
3.00	100%	100%	99%	98%	96%	94%	91%	86%	81%	38%	287.82
5.00	100%	100%	100%	100%	100%	99%	99%	98%	97%	81%	403.94

PROBABILITY THAT NO MORE THAN 10% OF AIRCRAFT ARE NMCS AT END OF MONTH 30

TABLE 2

V/M USED FOR STOCKAGE	VARIANCE-TO-MEAN RATIO USED TO EVALUATE PERFORMANCE										STOCKAGE COST $ (MILLION)
	1.00	1.25	1.50	1.75	2.00	2.25	2.50	2.75	3.00	5.00	
1.00	8.9%	10.3%	11.7%	13.0%	14.3%	15.5%	16.7%	17.9%	19.1%	27.7%	155.15
1.25	7.4%	8.7%	10.0%	11.3%	12.5%	13.7%	14.9%	16.1%	17.2%	25.7%	172.14
1.50	5.9%	7.2%	8.5%	9.7%	10.9%	12.0%	13.2%	14.3%	15.4%	23.8%	189.13
1.75	4.8%	5.9%	7.1%	8.2%	9.3%	10.5%	11.6%	12.7%	13.8%	22.0%	206.33
2.00	3.8%	4.5%	5.8%	6.9%	8.0%	9.1%	10.1%	11.2%	12.2%	20.3%	223.00
2.25	2.9%	3.8%	4.7%	5.7%	6.7%	7.7%	8.8%	9.8%	10.8%	18.7%	239.61
2.50	2.3%	3.0%	3.8%	4.7%	5.6%	6.5%	7.5%	8.5%	9.4%	17.1%	256.00
2.75	1.7%	2.4%	3.1%	3.8%	4.6%	5.5%	6.4%	7.3%	8.2%	15.6%	272.06
3.00	1.4%	1.9%	2.5%	3.1%	3.8%	4.5%	5.3%	6.2%	7.0%	14.2%	287.82
5.00	.64%	.71%	.80%	.90%	1.0%	1.2%	1.4%	1.7%	2.0%	5.9%	403.94

EXPECTED PERCENT OF AIRCRAFT NMCS AT END OF MONTH 30

TABLE 1

stockage planning purposes. This is the value that determines, in this study, which items will be purchased and to what quantity. The column to the extreme right in each table is the cost of the stock, in millions of dollars, purchased at each estimated variance-to-mean ratio to achieve an 80 percent probability of having 10 percent or fewer NMCS aircraft at any base. Across the top of each table is the actual variance-to-mean ratio that was used in the performance evaluation of each of the ten stockage policies. For example, the cell in the upper right corner of each table reflects the results that would be obtained should the planner assume a variance-to-mean ratio of 1.0, and purchase stock using this variance-to-mean ratio, but the actual variance-to-mean ratio encountered was 5.0 in the operational environment.

Each cell in the matrix represents the intersection of the estimated variance-to-mean ratio and the actual. The statistic in each cell of Table 1 is the percentage of NMCS aircraft for all bases, at the end of month 30, assuming full cannibalization. Dyna-METRIC generated for each base a point estimate of the expected number of NMCS aircraft and a standard deviation of that point estimate. The statistic in each cell is the total of the point estimates divided by the total number of aircraft available at the end of month 30.

The statistic in each cell of Table 2 is the prob-

ability or the percent of time, that a specific situation will produce 10 percent or fewer NMCS aircraft under a full cannibalization scenario. For example, the cell mentioned before (1.0 versus 5.0) has a zero chance of producing 10 percent or fewer NMCS aircraft. This statistic is a weighted average of all eight bases. Each base had its own probability computed. The figure presented is the average of all eight locations based on the percentage of aircraft at that location at the end of month 30.

Figure 3 shows the relationship of the estimated variance-to-mean ratio to the cost of the stockage policy. It was developed from the tables to aid in answering research question number one.

Figures 4 and 5 are contour graphs developed from Tables 1 and 2 respectively. These will aid in the analysis of research question number two. Figure 4 shows what happens to the point estimate of the expected percent of NMCS aircraft (Y axis) when the actual variance-to-mean ratio (X axis) is not the same as the estimated (contour lines). Figure 5, on the other hand, displays how the probability of achieving a 10 percent NMCS rate (Y axis) is affected by experiencing an actual variance-to-mean ratio (X axis) which differs from a given estimated variance-to-mean ratio (contour lines).

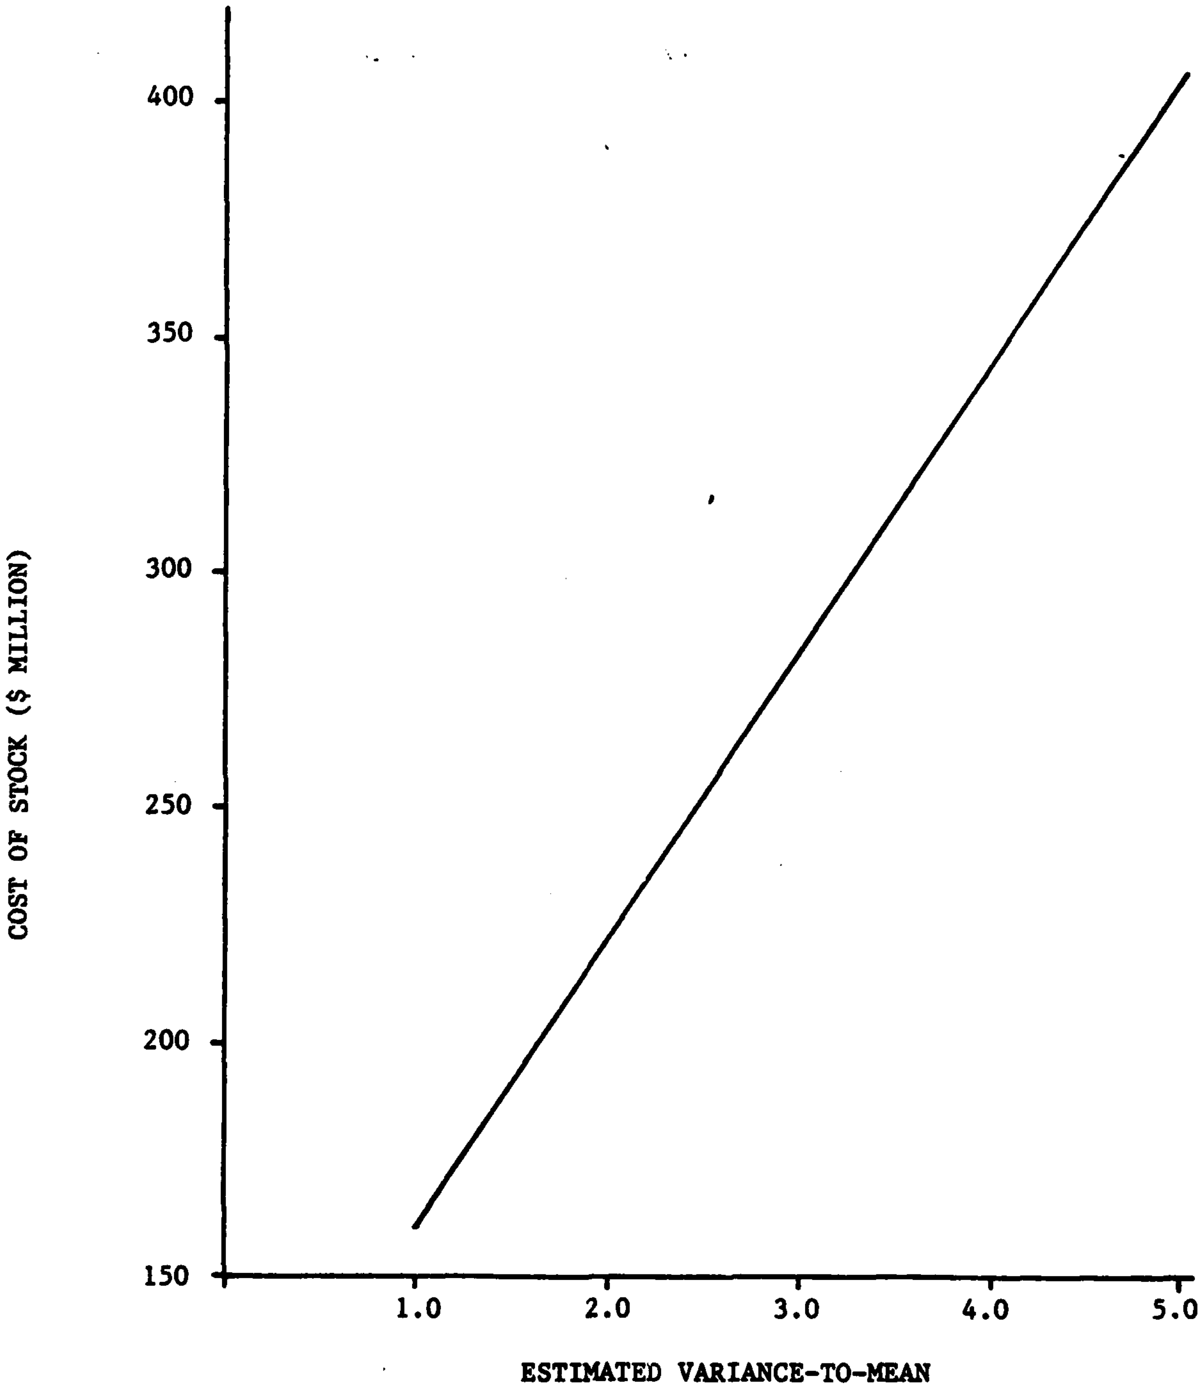

Figure 3. Effect of Variance-to-Mean On Cost Of Stock

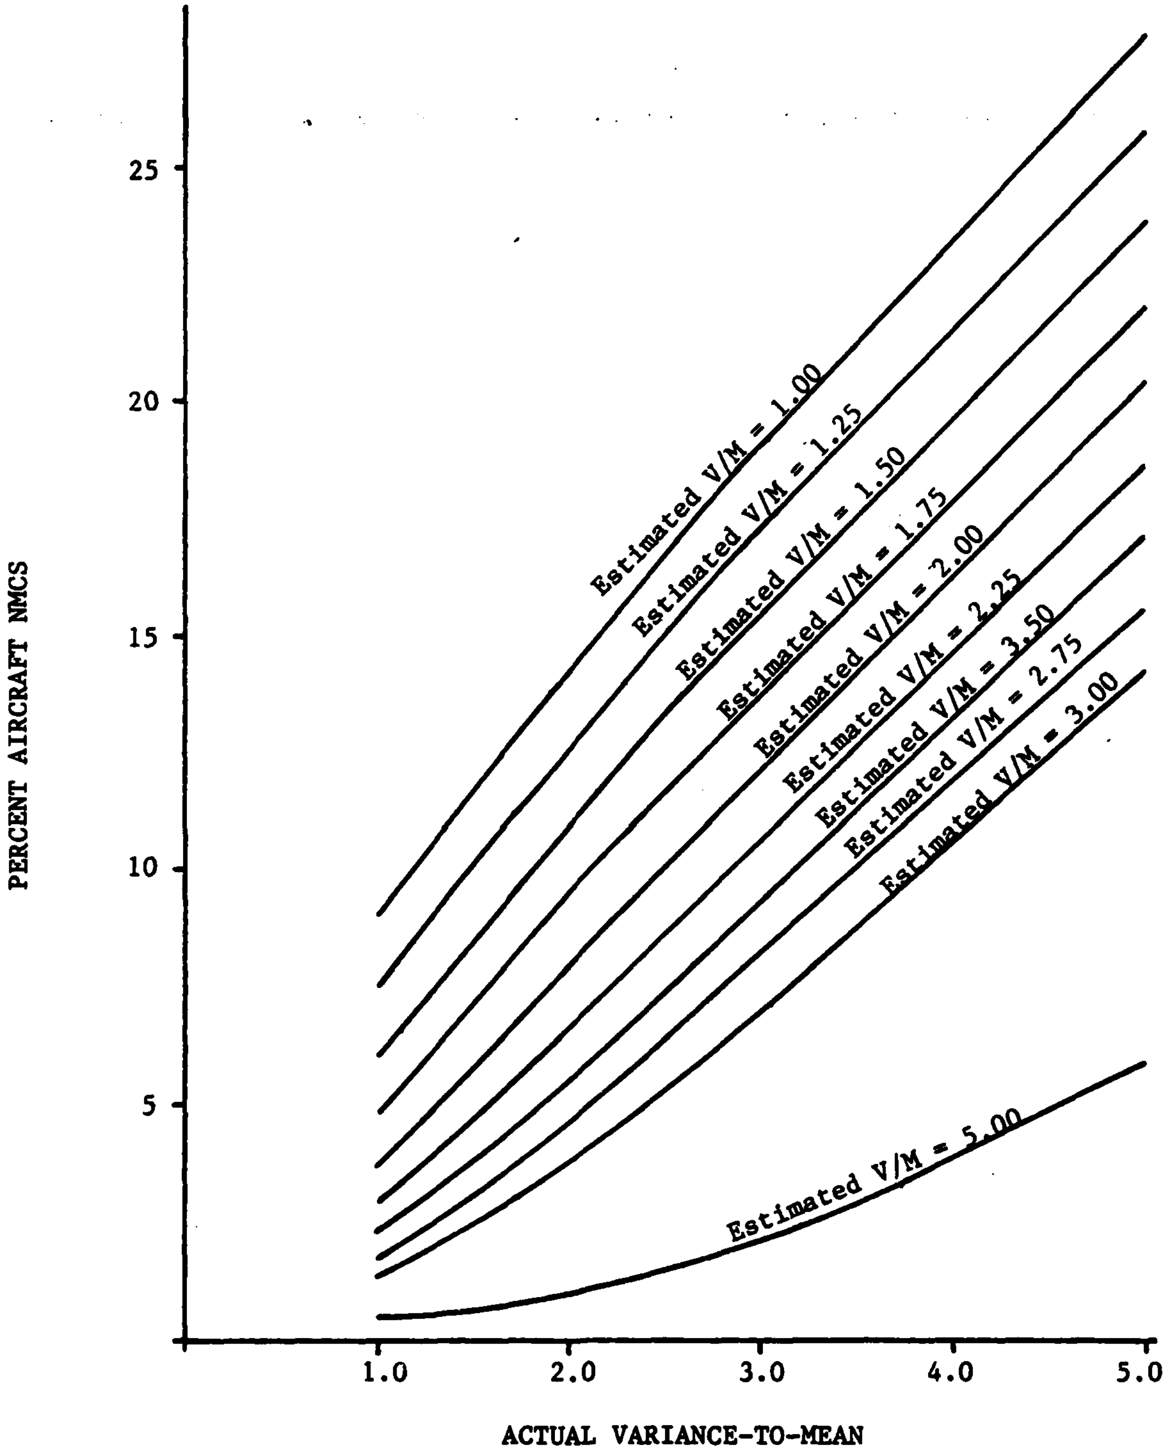

Figure 4. Percent Aircraft NMCS

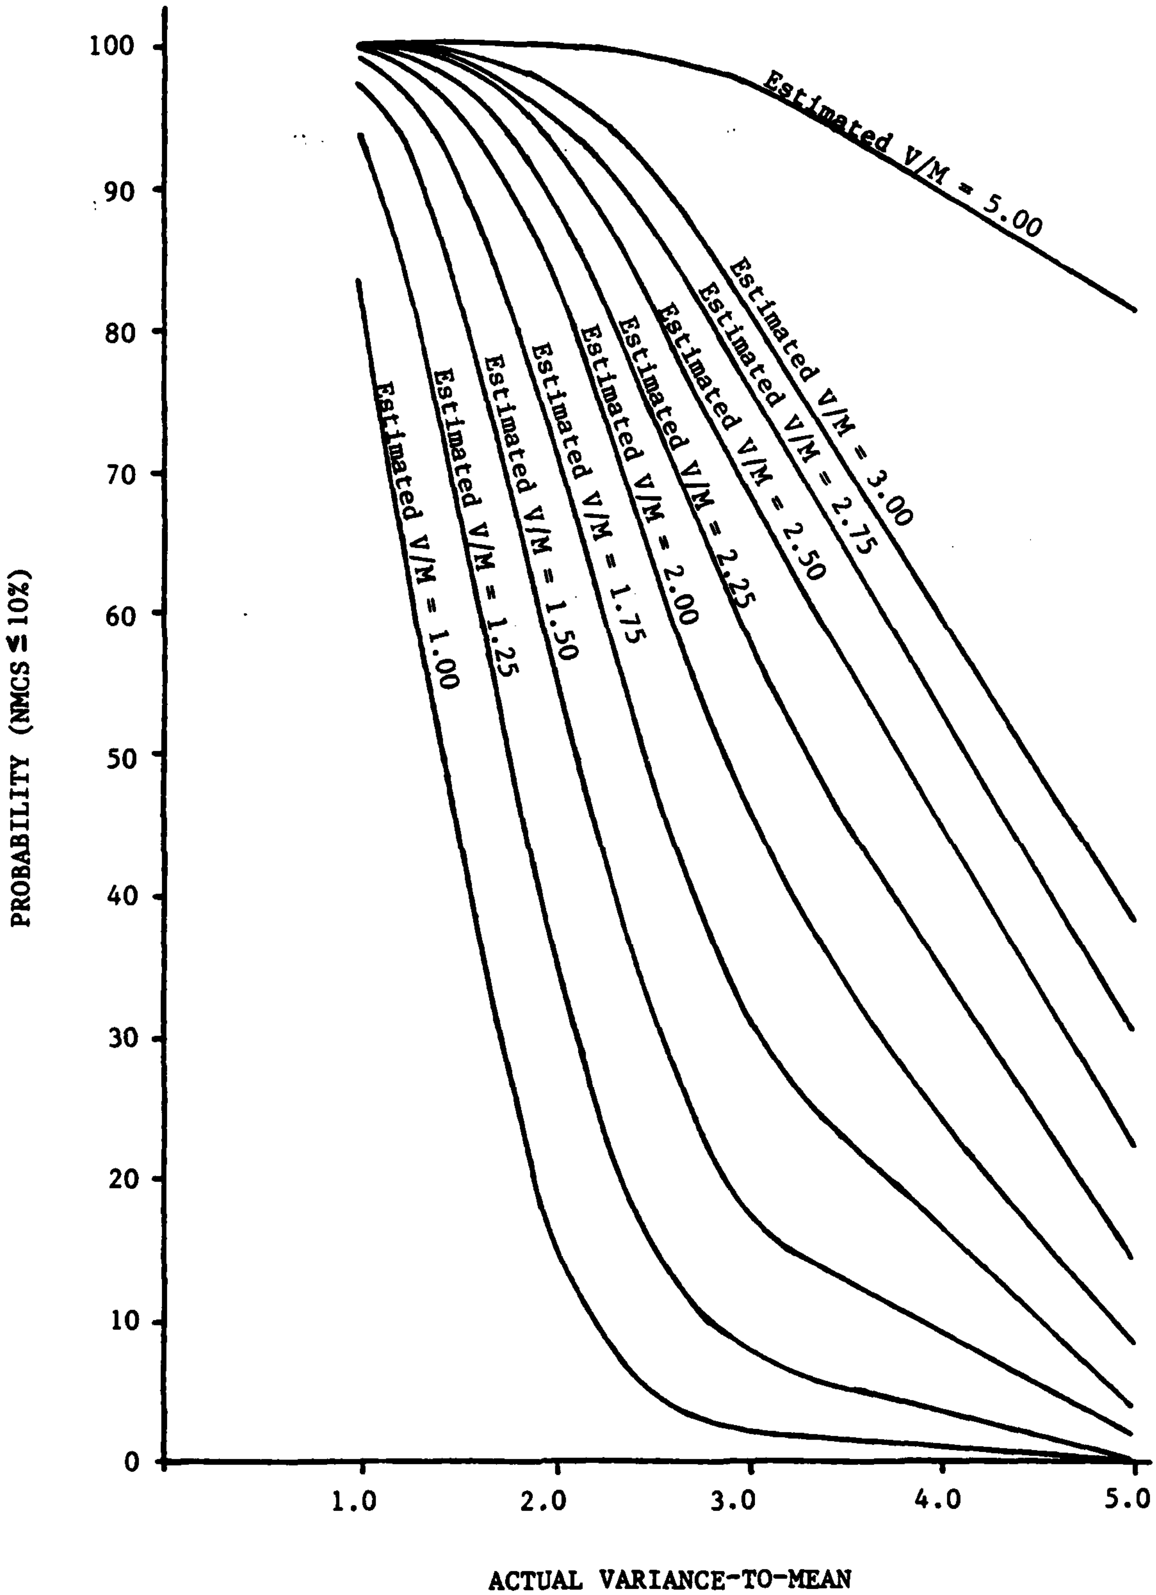

Figure 5. Probability of Achieving Performance Target

Analysis

A question arises when looking at the figures in each cell along the upper left to lower right diagonal of Table 1. This is the diagonal where the estimated and actual variance-to-mean ratios are equal. Intuitively one might think they should all be alike. At 1.0 versus 1.0, the expected NMCS percent is 8.9 percent. But that value tapers significantly to 5.9 percent at 5.0 versus 5.0. This can be attributed to the variability of the demand distribution. The expected NMCS percent or mean percentage can be close to the 10 percent NMCS goal when demand variability is low (e.g., variance-to-mean ratio = 1.0) to achieve the 80 percent probability of meeting the 10 percent target. However, as demand variability increases, the expected NMCS percent must be made lower, because one has to go out further to the right in the demand distribution to reach the 80 percent probability density point of a 10 percent NMCS rate. The bell-shaped distributions in Figure 6, although not totally accurate with respect to the exact distribution of NMCS aircraft, serve to illustrate this point.

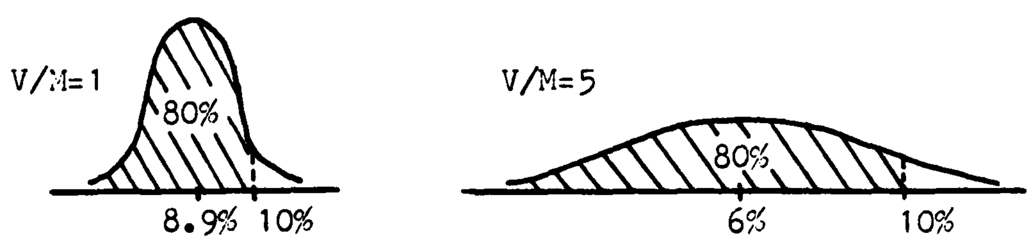

Figure 6. NMCS Percentage Distributions: V/M=1 and V/M=5

When the estimated and actual variance-to-mean ratios are equal, as viewed on the upper left to lower right diagonal of Table 2, one might intuitively expect that the management goal of achieving an 80 percent probability of 10 percent or fewer NMCS aircraft would be hit exactly. As observed in the figures of each of these diagonal cells, the goal was exceeded. The probabilities of achieving the 10 percent would have been achieved if the scenario was totally static. However, aircraft levels and flying hour programs changed periodically throughout the scenario time frame. The Dyna-METRIC model did just as it should have done. It computed required stock levels to achieve the 80 percent target throughout the scenario. Some of the base sortie rates peaked at month 22 and tapered off to month 30. The model bought stock to cover the peak program at month 22 and, of course, had more than enough stock to meet the 80 percent goal at month 30, resulting in performance greater than 80 percent.

Figure 3 is a graph of the relationship between the estimated variance-to-mean ratios and cost in millions of dollars. As can be seen, for all practical purposes, the relationship is linear. The slope of the line indicates that for a unit increase in the variance-to-mean ratio, costs go up by 62.20 million dollars from a baseline of $155.15 million. That translates to an approximate 40 percent increase in costs for each unit increase of the

variance-to-mean ratio over the baseline of 1.0 A stock computation done with the variance-to-mean ratios set at 5.0 will result in a 160 percent increase in stockage investment cost over a computation done at 1.0.

An analysis of Figure 4, which shows the effect of estimated versus actual variance-to-mean ratios on the point estimate of the NMCS percentage rate, shows some sensitivity of the NMCS percentage to changes in the variance-to-mean ratio. An analysis of the two extreme contour lines (estimated variance-to-mean = 1.0 and estimated variance-to-mean = 5.0) will provide the general trends. When stock is bought at a variance-to-mean of 1.0 in this scenario it will cost 155.15 million dollars. However, when one assumes this low variability of demand and it does not turn out to be that low in actuality, performance suffers. It can be seen that on the estimated variance-to-mean = 1 contour line that when the actual variance-to-mean ratio varies from 1.0 to 5.0, the expected NMCS percent goes from 8.9 percent to 27.7 percent of the entire fleet of aircraft. On the other hand, when stock is bought at a variance-to mean ratio of 5.0 in this scenario, it will cost $403.94 million. However, when variability of demand is not as estimated the performance does not suffer by the same magnitude. In fact, when the actual variance-to-mean ratios vary from 1.0 to 5.0 (estimated variance-to-mean = 5.0 contour line), NMCS percentage goes from .64 percent

to 6 percent. The extremes indicate relatively low cost at low variance-to-mean ratio, but high risk performance when the variance-to-mean ratio is not as expected, and relatively high cost at high variance-to-mean ratio, with resulting low risk performance when the variance-to-mean ratio is not as expected.

Figure 5 shows how the probability of achieving a 10 percent NMCS rate is affected by experiencing an actual variance-to-mean ratio which differs from an estimated variance-to-mean ratio. Again, the two extreme contour lines will be analyzed. Trend patterns between these extremes are easily recognizable. The bottom contour (estimated variance-to-mean = 1.0) shows that the probability of achieving the desired NMCS rate of 10 percent drops rapidly at a decreasing rate. Even when the actual variance-to-mean ratio was only one above the estimated, or 2.0, the probability of achieving the target is only 15 percent. On the other hand, when the top contour line (estimated variance-to-mean = 5.0) is analyzed between actual variance-to-mean ratios from 1.0 to 5.0, the performance goal of an 80 percent probability of achieving a 10 percent NMCS rate was far exceeded even at a variance-to-mean ratio of 3.0, where there is still a 97 percent probability of achieving the goal. It must be kept in mind, however, that as estimated variance-to-mean ratios increase, so does the cost of stock. Although high estimated variance-to-mean

ratios (e.g., 5.0) will assure performance for actual variance-to-mean ratios less than 5.0, there is a cost: in this scenario it is $403.94 million. A low estimated variance-to-mean ratio (e.g., 1.0) is relatively inexpensive ($155.15 million), but if the real variance-to-mean ratio is higher, the probability of achieving a NMCS goal will drop at varying rates depending on which contour line stockage computations have been done.

Summary

Stockage costs increase as the estimated variance-to-mean ratios are increased. In this scenario, with a variance-to-mean ratio of 1.0 as the baseline, stock costs increased approximately 40 percent for each unit increase in the variance-to-mean ratio. While costs were $155.15 million at a variance-to-mean ratio of 1.0, they were $403.94 million at 5.0.

There are cost/performance tradeoffs when the actual variance-to-mean ratios experienced in an operational scenario are not as estimated during stockage computations. Stock computations done at a lower variance-to-mean ratio than actually experienced will be less expensive than if the actual variance-to-mean ratio had been known and used in the stockage computation. However, performance, in terms of the actual percent of NMCS aircraft and to a greater degree, the probability of achieving a 10 percent NMCS rate,

will not be as good as expected. Stock computation at a variance-to-mean ratio of 1.0 and performance at 5.0 gives an expected NMCS percentage of 27.7 percent when a NMCS percent of 8.9 would have resulted had the actual and estimated variance-to-mean ratios both been equal to 1.0. An estimated and actual variance-to-mean ratio of 1.0 would achieve an 83 percent probability, but, when the actual variance-to-mean is 5.0, the probability drops to 0 percent. Stock computations done at a higher variance-to-mean ratio than actually experienced in the field will help assure that performance targets will be met. Stock computations done at an estimated variance-to-mean ratio of 5.0, with actual variance-to-mean ratio at 1.0 resulted in a 100 percent probability of achieving the 10 percent NMCS rate. This is overperformance, since the target was 80 percent. This overperformance comes at a cost. Stock computations in this instance result in more dollars spent than needed to simply achieve the 80 percent goal.

CHAPTER V

CONCLUSIONS AND RECOMMENDATIONS

Summary of Research Effort

Background. This research examines the sensitivity of certain recoverable item inventory model outputs to changes in a specific input parameter. That input parameter, the variance-to-mean ratio, is the measure of dispersion (or uncertainty) of item demand about the mean demand. One of the model outputs is dollars of stock necessary to achieve desired performance goals (the probability of having at least a desired number of available aircraft). The other output is performance (probability of having at least a desired number of aircraft available) given a certain stockage policy (i.e., quantities and mix of spare parts).

The variance-to-mean ratio is necessary to express the uncertainty of demand and is used in all contemporary recoverable item inventory models in some form. Beginning with the METRIC model developed by Sherbrooke, all algorithms use the variance-to-mean ratio in basically the same way. Sherbrooke and others found the compound Poisson to better describe the demand distribution than a simple Poisson and used the negative binomial distribution

to approximate it. In spite of the fact that Sherbrooke found METRIC to be extremely sensitive to the variance-to-mean ratio, there exists many different viewpoints on the importance of determining the correct variance-to-mean ratio. Twenty years after the pioneering research, there is still no agreed upon method of determining the correct variance-to-mean ratio. Some Air Force models such as WRSK/BLSS and the Combat Support Capability Management System still use the simple Poisson with a variance-to-mean ratio of 1.0, even though it might be too simple to describe reality. The rationale for using the simple Poisson seems to result from a lack of a clear understanding of the sensitivity of the models to the variance-to-mean ratio coupled with the traditional belief that failure rates are simple Poisson. Other models such as the AFLC Recoverable Items Requirements Computation System (D041A), and the AFLC Central Leveling System (D028A) recognize the fact that the variance-to-mean ratio can vary by item and use a power function based on mean demand to compute the variance-to-mean ratio for each item. The LMI Availability model increases the variance-to-mean ratio for outyear projections to cover for uncertainty of demand. The WARS model uses the square root of the number of bases, not to exceed 5.00 as the estimated variance-to-mean ratio.

With the current investment in recoverable spares of 2.5 billion dollars, being sure that money is being

wisely spent to achieve desired warmaking capability is of utmost importance.

Problem Statement and Research Questions. There is much concern over how the variance-to-mean ratio is determined and what it should be for any given item of stock. Logisticians, for the most part, feel that model outputs are sensitive to the variance-to-mean, but don't know the magnitude of sensitivity. The objective of this research is to give logisticians some insight into just how their decisions in establishing the variance-to-mean ratio, for each recoverable item, affects the cost and performance outputs of the models they are working with. The criticality of their decisions in this regard depend upon the sensitivity of outputs to the variance-to-mean ratio.

There are two research questions. The first asks whether the contemporary recoverable item model output of cost of stock required to achieve a desired performance parameter is sensitive to changes in the variance-to-mean ratio. The second asks whether, when computing the performance (i.e., probability of achieving a desired NMCS level) of a group of aircraft, given a likely scenario and a fixed quantity and mix of spare parts, the model output of mission performance is sensitive to changes in the variance-to-mean ratio.

Methodology. Dyna-METRIC, a RAND developed recoverable item inventory model, and a validated F-16 data base

obtained from Ogden Air Logistics Center were used to answer both research questions. A ten by ten matrix was used for the experimental design. Ten stock computations were made with ten individual variance-to-mean ratios between 1.0 and 5.0. The effort was concentrated in the 1.0 to 3.0 range where 80 percent of AFLC's recoverable items are. Those stock computations, and the associated cost of stock, addressed research question number one. For each individual stock computation, ten individual scenario performance computations were made to answer research question number two. All parameters except those mentioned above were held constant. All variance-to-mean ratios for all items were changed to the same value for each model run.

Conclusions

Research Question Number One. The dollars of stock necessary to achieve certain scenario performance objectives are sensitive to the variance-to-mean ratio. As the results show, there is an approximate 40 percent increase in investment required for each unit increase in the variance-to-mean over 1.0. Assuming that the Air Force required $2.0 billion of spares at a variance-to-mean ratio of 1.0, it would require $3.6 billion at a variance-to-mean ratio of 3.0. The models are, therefore, investment sensitive to the variance-to-mean ratio.

Research Question Number Two. The model outputs of mission performance are very sensitive to changes in the variance-to-mean ratio. When the actual variance-to-mean ratio was 100 percent greater than the estimated, the probability of achieving the NMCS goal of 10 percent or fewer aircraft decreased by approximately 60 percentage points. There is virtually no chance (i.e., 2 percent probability) that the 10 percent NMCS goal would be reached if the stock computation were based on a variance-to-mean ratio of 1.0 and the actual scenario performance experienced was at 3.0. Also, the expected number of NMCS aircraft increased by anywhere from 50 to 100 percent, depending on whether estimate versus actuals were 1.0 versus 2.0, 2.5 versus 5.0, or any combination in between where the actual is two times the estimated.

Implications for Recoverable Item Inventory Managers

By stocking according to variance-to-mean ratios which differ from what would be experienced in the field, the Air Force could be unknowingly making cost/performance tradeoffs with far reaching consequences. If the variance-to-mean ratio used for stock computations is estimated to be lower than it really is, money will be saved. It could be a lot of money, money that could be used for other worthwhile programs. But there is also a hidden cost which must be paid for this dollar savings. That cost is the

reduced combat capability that would be experienced. There would be a false sense of security that the stock levels computed at unrealistically low variance-to-mean ratios would support the desired level of aircraft availability. Conversely, if the estimated variance-to-mean ratio is overstated during stock computation, the combat unit will have a much higher probability of achieving its goals due to the excess stock that will be purchased. However, this over-performance comes at a price. The price this time is that more money than necessary was spent to achieve the specified level of performance. This is money that could be spent on other necessary programs.

Also, there is a very real danger in using these contemporary recoverable item inventory models in evaluating currently existing stockage policies for any weapon system. That danger is incorrectly determining what variance-to-mean ratios to use for the stock items. If they are set too low, the performance output could indicate a false sense of security. Managers might even think they are overstocked and cut back on stock or defer future programmed buys. When it came to actually exercising that stock in a war, the managers might find that their action to trim stocks caused under-performance and loss of a battle or a war. On the other hand, if the variance-to-mean ratios are set too high, performance outputs might lead Air Force managers to think they are understocked. They might

ask and get funding from Congress to buy the extra stock they feel they need but which is not in fact necessary.

This research has implications for all weapon systems, not just the F-16. Further studies using data bases from other weapon systems should substantiate this.

Recommendations

The variance-to-mean ratio should be set as accurately as possible to ensure that performance goals are met, at the least possible cost. The credibility of the Air Force, before Congress, can only be enhanced by requesting no more funds than necessary and performing exactly as predicted. The authors applaud the effort currently under way at AFLC. A contract has been recently let to determine what the variance-to-mean ratios should be for recoverable items at AFLC. It is a limited effort, but is at least a start in the right direction.

Suggestions for Further Research

This study dealt only with the F-16 weapon system. To insure that generalizations of this research are accurate, other weapon system data bases should be studied with Dyna-METRIC to see if they produce comparable results.

Also, the F-16 and other data bases should be studied using other contemporary models (e.g., WARS, when it becomes available) to see if comparable results are produced.

A study should be conducted to see what effects random assignment of variance-to-mean ratios have on performance and cost as compared to the uni-directionally biased variance-to-mean ratios of this study.

It would be informative to investigate the effect on performance and cost of changing other variables in addition to variance-to-mean ratios.

This research was done using Dyna-METRIC under a full cannibalization assumption. In the future, RAND Corporation should have the capability built into Dyna-METRIC to use the no cannibalization assumption. This research should be repeated using that assumption. This would give the other extreme of the cost/performance envelope. Reality, more than likely, is somewhere in between the extremes.

Using a dynamic cost constrained inventory model, such as WARS, in the stock computation mode, it would be most interesting to continue Sprung's research and find out what the spare parts mix (i.e., quantities and types) would be for different estimated variance-to-mean ratios. If it turns out that the Air Force would buy the same mix of parts for variance-to-mean ratio = 1 as it does for variance-to-mean = 3, given likely funding levels, then it does not matter for actual stock computation purposes (other than funding requests) what the variance-to-mean is set to.

A Final Note

If the Air Force is to be mean and lean, it must use the correct variance-to-mean.

APPENDIX A

OGDEN ALC F-16 DATA BASE

The following tables and discussion are intended to show the realism of the data base used. The data base was an up-to-date version as of the second quarter, FY 82. Ogden uses this data base with Dyna-METRIC.

As mentioned before, the data is scaled by a factor of 30, so that one day in the data base equals one month. All variables and parameters have been scaled appropriately.

There are three USAFE and five PACAF bases in this peacetime scenario. A depot is modeled in the place of the CIRF. Ogden uses Dyna-METRIC against their data base to determine the performance of USAFE/PACAF bases given certain stockage for each base and the depot at specific points in time.

The overall structure of the data includes: 1) random exponential repair times, and, 2) no base and depot administrative delays.

The depot scenario specifications (since depot is substituted for the CIRF) are:

1. Intermediate Level Maintenance (ILM). Deployment (in this case depot maintenance) for remove, replace, repair (RRR) items occurs at month 0.
2. ILM setup period is at month 0.
3. Beginning day of CONUS resupply is at month 0.
4. ILM deployment for remove, replace (RR) items is at month 0.
5. ILM setup for RR items is at month 0.

AD-A124 403 THE EFFECTS OF ITEM USAGE VARIATION ON INVENTORY STOCKAGE MODELS(U) AIR FORCE INST OF TECH WRIGHT-PATTERSON AFB OH SCHOOL OF SYST.. 2/2

UNCLASSIFIED J H KUTZKE ET AL. SEP 82 AFIT-LSSR-77-82 F/G 15/5 NL

END

FILMED

DTIC

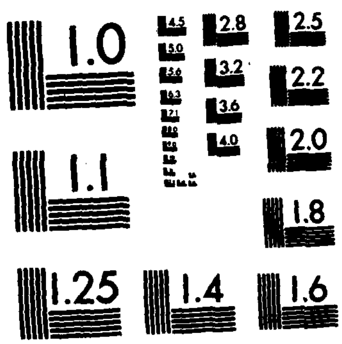

MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

6. CONUS peacetime pipeline empties from month 0.

7. Full SRU cannibalization.

8. SRU repair capability is available on month 0.

9. Peacetime SRU resupply time from next higher echelon of supply is two months.

10. SRU resupply is available at month 0.

11. First day of forward transportation cutoff from next higher echelon of supply to depot is at month 0.

12. Duration of cutoff from next higher echelon of supply is 0 months.

The base scenario specifications are identical to the depot's, except that where the words depot and next higher echelon of supply appear, the words base and depot respectively should appear. Also, there are two additional parameters:

1. Transportation time from depot to base

   PACAF=.5 month   USAFE=.17 months

2. Transportation time from base to depot

   PACAF=.77 months   USAFE=.30 months

The aircraft level specifications are based on when F-16 squadrons will be activated and when their expected numbers will be increasing/decreasing at each base.

Numbers of Aircraft by Base

Base	Month 0	Month 4	Month 7	Month 13	Month 16	Month30
1	72	72	72	72	72	72
2	48	48	48	48	48	48
3	0	0	0	0	72	72
4	96	96	96	96	96	96
5	86	86	86	86	86	86
6	72	76	76	70	72	72
7	0	0	48	48	48	48
8	0	0	0	72	72	72

Sortie rate specifications were based on projections from prior scheduled versus actual performance of F-16 or similar aircraft at each base. Months on which sortie rates change are based on changes in programmed flying hours.

Sorties Per Aircraft Per Month By Base

Base	Mo 0	Mo 4	Mo 7	Mo 10	Mo 13	Mo 16	Mo 19	Mo 22	Mo 25
1	0.0	12.7	14.2	15.3	15.5	16.7	17.9	17.9	17.9
2	12.2	15.3	15.5	15.5	16.7	17.9	18.1	20.2	22.9
3	0.0	0.0	0.0	0.0	0.0	16.7	16.9	16.9	16.9
4	16.7	16.8	18.1	19.1	18.3	18.3	18.9	20.0	19.2
5	16.7	16.8	18.1	19.1	18.3	18.3	18.9	20.0	19.2
6	16.7	16.8	18.1	19.1	18.3	18.3	18.9	20.0	19.2
7	0.0	0.0	18.1	19.1	18.3	18.3	18.9	20.0	19.2
8	0.0	0.0	0.0	0.0	18.3	18.3	18.9	20.0	19.2

Flying hours per base are as programmed and are static throughout the scenario.

Flying Hours/Sortie/Aircraft

Base	Flying Hours
1	1.2
2	1.2
3	1.2
4	1.4
5	1.4
6	1.4
7	1.4
8	1.4

There were 282 LRUs and 236 SRUs in the scenario. Exhibit 1 is a representative sample of 10 of these LRUs and 5 SRUs and their parameters. Parameters are based on actual experience. Please note that the NRTS rate at the depot of 1.000 means that the part is repaired at the depot with the peacetime order and ship time being the average repair time at the depot. Also, all items are removed, repaired, and replaced (RRR). The linearity factor for LRUs is 1.000 because of peacetime scenario. If this were wartime, the linearity factor could be changed to accelerate/decelerate parts usage rates. Level of repair for SRUs is at the base or depot.

All spare parts were used at all bases in this scenario.

The following is an example of an indentured relationship used in the scenario:

Indentures

Part No. LRU	Part No. SRU	QPA
1270010932174WF	1270010956768WF	1
	1270010600697WF	1
	1270010615082WF	1
	1270010778077WF	1
	1270010778197WF	1
	1270010774184WF	1

The stock levels in the original OOALC data base were not used. As mentioned in the methodology section, stock levels for this research were computed by Dyna-METRIC.

LRUs

Part No.	Demand per Fl Hr	NRTS Rate Base	NRTS Rate Depot	Base Repair Time(Mo)	Unit Cost	Order & Ship Time(Mo)
1560010440179WF	.00001	.000	1.000	0.200	$ 2285	1.633
1560011067173WF	.00011	.990	1.000	0.200	2240	1.633
1630010454508	.00055	.750	1.000	0.200	3990	3.333
5820010376772	.00145	.480	1.000	0.100	662	2.467
5895005391911	.00377	.300	1.000	0.200	17,922	1.000
1280011096916WF	.00474	.230	1.000	0.133	54,953	1.367
5841010963945WF	.00522	.270	1.000	0.167	26,980	1.133
1270010932256WF	.00715	.840	1.000	0.133	76,119	1.933
6605010876645	.00931	.510	1.000	0.167	198,961	1.933
1630010389239	.02000	.150	1.000	0.267	2072	1.300

SRUs

Part No.	Demand per Fl Hr	NRTS Rate Base	NRTS Rate Depot	Base Repair Time(Mo)	Unit Cost	Order & Ship Time(Mo)
1620010613240	.00001	.000	.000	0.200	1450	2.867
6615010720097WF	.00060	.000	.000	0.200	2227	1.700
1270010967120WF	.00038	.000	.000	0.200	44,180	1.700
5895005642041	.00108	.940	.940	0.100	3729	1.033
1270010830398WF	.00196	.000	.000	0.200	56,571	1.700

LRU/SRU Description Data

Exhibit 1

SELECTED BIBLIOGRAPHY

A. REFERENCES CITED

1. Brooks, Robin B. S., Colleen A. Gillen, and John Y. Lu. Alternative Measures of Supply Performance: Fills, Backorders, Operational Rate, and NORS. The RAND Corporation, RM-6094-RR, Santa Monica CA, August 1969.

2. Crawford, G. B. Palm's Theorem for Nonstationary Processes. The RAND Corporation, R-2750-RC, Santa Monica CA, October 1981.

3. Department of Defense. Contract SD-271 with Logistics Management Institute. Washington, 14 July 1971.

4. Duke, First Lieutenant James A., USAF, and First Lieutenant Kenneth W. Elmore, USAF. "An Empirical Investigation of the Effects of Inventory Stockage Models for Recoverable Items on Weapon System Availability." Unpublished Master's thesis. LSSR 14-81, AFIT/LSH, Wright-Patterson AFB OH, June 1981.

5. Estey, Arthur S., and Ronald L. Kaufman. "Multi-Item Inventory System Policies Using Statistical Estimates: Negative Binomial Demands (Variance/Mean = 9)." Unpublished Technical Report No. 3, Mathematical & Information Science Division, Offices of Naval Research (Code 434), Arlington VA, September 1975. AD A015810.

6. Faucett, W. M., and R. D. Gilbert. Characteristics of Demand Distributions for Aircraft Spare Parts. General Dynamics Document RDP 414-61-500, Fort Worth TX, 10 November 1966.

7. Feeney, G. J., and C. C. Sherbrooke. The (s-1,s) Inventory Policy Under Compound Poisson Demand: A Theory of Recoverable Item Stockage. The RAND Corporation, RM-4176-PR, Santa Monica CA, September 1964.

8. ________. A System Approach to Base Stockage of Recoverable Items. The RAND Corporation, RM-4720-PR, Santa Monica CA, December 1965. AD 627644.

9. Feeney, G. J., J. W. Peterson, and C. C. Sherbrooke. An Aggregate Base Stockage Policy for Recoverable Spare Parts. The RAND Corporation, RM-3644-PR, Santa Monica CA, June 1963. AD 408943.

10. Galliher, Herbert P., and Richard C. Wilson. "Aircraft Engines: Demand Forecasting and Inventory Redistribution." Unpublished research report No. ARL 75-0008, Applied Mathematics Research Lab/LB, Aerospace Research Laboratories (AFSC), Wright-Patterson AFB OH, February 1975. AD A011595.

11. Hales, Kenneth. Operations Research Analyst, OOALC/MMM-2, Hill AFB UT. Telephone interview. 3 August 1982.

12. Hill, John. Supervisory Operations Research Analyst, Management Sciences, DCS/Plans and Programs, HQ AFLC, Wright-Patterson AFB OH. Personal interview. 3 August 1982.

13. Hillestad, R. J. Dyna-METRIC: Dynamic Multi-Echelon Technique for Recoverable Item Control. The RAND Corporation, WD-911-AF, Santa Monica CA, September 1981.

14. Hillestad, R. J., and M. J. Carrillo. Models and Techniques for Recoverable Item Stockage When Demand and the Repair Process Are Nonstationary--Part I: Performance Measurement. The RAND Corporation, N-1482-AF, Santa Monica CA, May 1980.

15. Isaacson, Karen E. Resident Consultant, System Science Department, The RAND Corporation, Santa Monica CA. Telephone interview. 21 July 1982.

16. Johnson, Captain David R., USAF, and Judith M. McCoy, GS-12. "An Analysis of the Demand Distribution for Selected Inertial Measurement Units in the Air Force Inventory." Unpublished Master's thesis. LSSR 5-78A, AFIT/LS, Wright-Patterson AFB OH, June 1978. AD A059287.

17. Management Sciences, DCS/Plans and Programs, HQ AFLC. "Wartime Assessment and Requirement System Summary Report." Wright-Patterson AFB OH, April 1981.

18. Muckstadt, J. A. Comparative Adequacy of Steady-State Versus Dynamic Models for Calculating Stockage Requirements. The RAND Corporation, R-2636-AF, Santa Monica CA, November 1980.

19. Muckstadt, John A. "A Model For a Multi-Item, Multi-Echelon, Multi-Indenture Inventory System," Management Science, Vol 20, No. 4 (December 1973), pp. 472-481.

20. Niklas, Michael R. Operations Research Analyst, Management Sciences, DCS/Plans and Programs, HQ AFLC, Wright-Patterson AFB OH. Personal interviews conducted intermittently from 17 May 1982 to 13 August 1982.

21. Presutti, Victor. Deputy Director, Management Science, DCS/Plans and Programs, HQ AFLC, Wright-Patterson AFB OH. Personal interview 3 August 1982.

22. Pyles, Raymond. The Dyna-METRIC Readiness Assessment Model: Motivation, Capabilities, and Uses. The RAND Corporation, WD-892-AF, Santa Monica CA, February 1982.

23. Pyles, Raymond, and Lieutenant Colonel Robert Tripp, USAF. Measuring and Managing Readiness: The Concept and Design of the Combat Support Capability Management System. The RAND Corporation, N-1840-AF, Santa Monica CA, April 1982.

24. Sherbrooke, Craig C. METRIC: A Multi-Echelon Technique for Recoverable Item Control. The RAND Corporation, RM-5078-PR, Santa Monica CA, November 1966.

25. ________. "METRIC: A Multi-Echelon Technique for Recoverable Item Control." Operations Research, Vol. 16 (1968), pp. 122-141.

26. Sprung, James G. "A Study of Variance-To-Mean in the OFM Requirement Algorithm." Unpublished working paper No. 44, Computer Science Division, Office of the DCS/Comptroller, HQ AFLC, Wright-Patterson AFB OH, May 1972.

27. Ste Marie, Paul. Deputy Director, Systems and Resource Management, HQ AFLC/LOM. Letter, subject: Study Topics for XRS In-House Effort, to HQ AFLC/XRS, 1 July 1981.

28. Stevens, R. J., and J. M. Hill. "Estimating the Variance-to-Mean Ratios for Recoverables in the ALS Marginal Analysis Algorithm." Unpublished working paper No. 49, Systems Study Branch, Offices of DCS/Comptroller, HQ AFLC, Wright-Patterson AFB OH, March 1973.

29. U. S. Department of the Air Force. Standard Base Supply System. AFM 67-1, Vol. II, Part Two, Chapter 11, pp. 3-13. Washington: Government Printing Office, 1980.

B. RELATED SOURCES

Feeney, G. J., and C. C. Sherbrooke. An Objective Bayes Approach For Inventory Decisions. The RAND Corporation, RM-4362-PR, Santa Monica CA, March 1965. AD 613470.

________. Systems Analysis and Supply Management. The RAND Corporation, RM-4054-PR, Santa Monica CA, April 1964. AD 437764.

END